

Date: Sun Nov 01, 1992 2:11 pm PST
Subject: manmade catastrophes

from Ed Ford (921101:1445) Bill Powers (921024.0830)

>I have nothing against trying to cope with manmade catastrophes, or
>any other kind, both before and after they occur. To do otherwise
>would be like giving up all attempts to control what matters to us,
>and I don't think that would happen even if I didn't agree that it's
>not a sensible thing to do. Coping with catastrophes and recovering
>from them demonstrate the basic human capacity to control and
>reorganize. But how well we do this depends on how well we understand
>what's going on. The first thing we have to understand is that people
>will continue to try to get what they want by any means that works,
>until they are dead. That's just a fact of nature. It is not changed
>by anything that either other people or the environment can do.....

I've just finished reading BANKRUPTCY 1995, The Coming Collapse of America and How To Stop It by Harry E. Figgie, Jr. After showing very convincingly how our annual increase in debt will lead to a complete collapse in our financial system by 1995, and I mean a real, total collapse, he suggests we deal with those who got us into this mess, namely the 536 members of Congress and the U.S. President. Then I read the above from Bill (I've been away in Seattle teaching PCT to school board members, administrators, and teachers from all over the state of Washington and a few adjoining states), including the statement "that people will continue to try to get what they want by any means that works, until they are dead." Kind of scary, to say the least. I do urge all of you to read the book, it tells where we are going in very clear, precise terms and it isn't pleasant. It does offer hope, but not for very long and, after having read Bill's remarks above, with only a slim chance of making a change before it's too late.

Gary, I appreciate your use of Freedom From Stress in your class. I am having printed a total revision of my Teaching Responsibility card and during the revision I came up with the name Teaching Responsible Thinking. It was so obvious. I have gotten rid of the concept of Responsible Behavior or Teaching Responsibility because, if PCT teaches anything, it teaches that we are educating others by teaching them how to think responsibly. Once I get the card, I'll send you some, which have the Love/Quality Time part on the reverse side, not as before on separate cards, cutting my costs in half. If anyone else would like some cards, let me know. They're free.

Ed Ford ATEDF@ASUVM.INRE.ASU.EDU
10209 N. 56th St., Scottsdale, Arizona 85253 Ph.602 991-4861

Date: Sun Nov 01, 1992 3:57 pm PST
Subject: linguistic units

[Avery Andrews 920211.1049]

A thought on linguistic structures (vaguely a propos of Sibun 9221026)

My belief is that most of the units people in linguistics talk about exist in some sense, but that the way in which they are normally thought about is probably seriously deficient and wrong-headed. Consider two of the better-established ones, 'noun-phrases', and 'clauses'. It seems to me that the S and NP nodes in linguistic structures could be interpreted as follows: when you're producing material under an NP, you're engaged in

`describing an entity', when you're producing material under an S, you're `describing a situation'. S and NP nodes & their corresponding rules appear to exist because there really are definite ways of proceeding to do these these things. The tree could be thought of as a sort of trace of the kinds of activities involved in producing or understanding the sentence.

But along with this element of truth, the tree-structure model people most assume these days has all kinds of much more dubious connotations, making straight-forward & plausible things seem difficult, & vice versa. For example, the existence of `locality constraints', limiting what parts of the tree can be simultaneously accessed by a single rule, seem to be rather deep and difficult, whereas they probably wouldn't be if there wasn't a picture to suggest that the whole structure exists in the head at the same time. For another example, consider the common `paratactic construction':

we walked up the hill, and then along a ridge, and then down into a valley, ...

In conventional tree-based thinking, you have to worry about whether to put an S over each little clause, and then one over the whole sequence, or whether there second one is under the S sitting over the first one, etc.:

 S S or or
???? S S S [stuff] S S

From a more process-based approach, a more straightforward answer comes to mind: whereas real subordination (e.g. the relation of the subject NP to the S it is the subject of) is a spawning relation between processes, the relation above is just one of chaining. E.g. there's no reason to think there's any master process that in controlling the transitions from one clause to the next. It's just that, after one has described one event, the natural thing to do next is to say what happened next.

One could build chaining into a tree-structure theory this way, perhaps:

S -> S -> S ...

(this is supposed to be a structure, not a rule), but the point is that there is nothing in the universe of tree-diagrams that makes this a particularly natural thing to do: the idea comes from thinking about how systems of processes might be organized.

So the general point is that while I think the notions found in linguistics are ways of talking about real things, I'm not committed to thinking that any of them are particularly good ways (like phlogiston can be seen as a way of talking about electrons, but everybody seems to agree that it's a pretty lousy one).

Avery.Andrews@anu.edu.au

Date: Mon Nov 02, 1992 10:01 am PST
Subject: description

[From Bill Powers (921102.0830)] Avery Andrews (921102) --

Avery, I think you are making a lot of sense: "It's just that, after one has described one event, the natural thing to do next is to say what happened next."

You're getting close to the way I think about the link between words and meanings (where, as you know, I define "meaning" as nonverbal perceptions). The key word here is _describe_. What is a description but a string of words that indicates some nonverbal happening? Just consider what underlies the words:

> we walked up the hill, and then along a ridge, and then down into a valley, ...

When you're describing something that already happened, you begin not with the words but with the memory of the walking tour, as it looked from your point of view, or perhaps from an imagined point of view as in looking at a map. Then you fashion words that would evoke images matching that memory in yourself. When others hear the words, they experience something that goes with them -- the images and so on. They aren't going to be the same images. I would hear the above phrases as describing a map-like or aerial scene, as seen from far off; if you had actually done that trip you might be hearing the same phrases as meaning the way the path looked to you from the ground, the slope of the hill rising ahead, the scene to the sides of the ridge, the view out over the valley as you descended into it.

In fashioning the phrases, you do so in the order that the experiences happened. You use the relationship-words that go with the perceive relationships: up the hill, but not up the ridge. You indicate the order of snapshots of the action with words designed to indicate succession: "and then." You describe what caused these things to happen: we walked.

I think description is the essence of the closed-loop relationship involved in creating utterances. It uses the same perceptual system that turns the words of others into perceptions of meanings, but it also involves the continuous creation of words such that they will evoke the desired meanings, at least in yourself. The "such that" part indicates the closure of the loop.

Best, Bill P.

Date: Mon Nov 02, 1992 10:06 am PST
Subject: Re: Meanings and sentences

[From Ray Allis (921102.0830)] [From Bill Powers (921027.0700)]

> >Variable aspects of sentences, for example (like meaning,
> >structure, inflection, etc) are controlled INPUTS, not generated
> >outputs.

> I want to make a modest push to make sure our linguists really truly
> get this point. They may actually get it, but I don't think we're
> hearing the result. Sentence construction is not construction of some
> object out there, or in some vague conceptual space; it's construction
> of an input, a perception. The mere fact that we know of a sentence
> shows that it's a perception. The same goes for grammar, for any
> regularity we PERCEIVE in language. The relationships between different
> levels of analysis of language are relationships among levels of
> perception, not levels of output production. We have to guess what the
> production processes are, because they're outputs and we don't perceive
> outputs. We perceive only their perceptual consequences: language is
> perception.

> [From Bill Powers (921025.1900)]
>
> P.S. Given three nouns A, B, and C, you can create the sentences
>
> The A B'd the C, and
> The B A'd the C.
>
> The captain hogged the ice-cream, and
> The hog captained the ice-cream.
>
> The only way to judge that one is allowable in ordinary discourse and
> the other is not is through the meanings of the words. That can't have
> anything to do with language itself, because the fact that hogs don't
> captain ice-cream is an accident of experience; given other
> experiences, we might see immediate sense in the second sentence but
> not the first, depending on what perceptions "hog" and "captain" are
> attached to, and what we can imagine hogs and captains doing. There's
> something about the way words are presented that connects to the way
> meanings are evoked in juxtaposition and sequence -- that's the input
> function above. But once the meanings are evoked, the rules that apply
> aren't those of language, but of perceived reality in general. We
> don't judge that hogs can't captain things because of language, but
> because of our experiences with hogs and captains.
>
> That must sound like utter nonsense.
> -----
> Best,
>
> Uncle Bill P.

"But once the meanings are evoked, the rules that apply
aren't those of language, but of perceived reality in general."

Gee, "Uncle Bill", I used to think I was the only one who thought in this utterly nonsensical way. The implications of this and some other related trains of thought are, in my experience, as invisible to "artificial intelligence" people as PCT is to psychologists. Maybe you'll be interested in this piece of something I haven't even *tried* to publish.

9.2 An example of analogy

An example of analogy (or more properly simile) is the phrase "Life is like a river". This sort of example has been criticized as being unrealistically difficult for computers. This phrase and many others are difficult even for humans. But that is the goal of AI, isn't it? The duplication and extension of human intellectual abilities?"

"Life is like a river." The "meaning" of this string of symbols has to do with time and change and irreversability, and with some things sensed and not so easily put into words, and this is the problem. Some or all of the similarity we perceive is non-linguistic, so that we cannot express these perceptions very well. (I am standing in the shallows, cold water pressing against my shins through the rubber of my hip boots, with a fly rod in my right hand.) I cannot express at all well in words my experience of "fitness" when I consider the similarity of the river and life. Much of my feeling is just that, feeling; non-verbal motor and sensory perception. Like the ability to ride a bicycle, experiences are not directly transferrable from mind to mind but must be acquired by each individual.

We can help each other, but we cannot give such experiences to others in the same way we can tell them how to multiply two numbers (i.e. manipulate deductive systems).

My concept and interpretation of this simile will certainly be different from yours, but sufficiently similar, we both believe, to enable communication. Our belief that we have communicated is an act of induction, based on our perceptions of behavior in others as being similar to our own behavior, and the hypothesis that similar actions are produced by similar motivations.

"Correctness" and "fit" have very different meanings here than in symbolic logic, no matter how "fuzzy" such logics are made. After all, what sense does it make to attempt to position symbols on scales of characteristics the symbols themselves do not have?

Similarity and analogy are relationships among representations (Section 7). The use of deduction (the treatment of the relationships among symbols) sacrifices the ability to perceive "likeness". You lose all the meaning!

Best to you, Ray

Date: Mon Nov 02, 1992 3:29 pm PST
Subject: Re: description

(penni sibun 921102.1200)

[From Bill Powers (921102.0830)]
Avery Andrews (921102) --

Avery, I think you are making a lot of sense: "It's just that, after one has described one event, the natural thing to do next is to say what happened next."

You're getting close to the way I think about the link between words and meanings (where, as you know, I define "meaning" as nonverbal perceptions). The key word here is describe. What is a description but a string of words that indicates some nonverbal happening?

well, i think you might enjoy my dissertation, which argued just these things (though in comp. sci. jargon rather than pct jargon). you and i disagree on the relationship bet. words and meanings, but i think that's actually an open question. what is more important is that we focus on the process.

I think description is the essence of the closed-loop relationship involved in creating utterances. It uses the same perceptual system that turns the words of others into perceptions of meanings, but it also involves the continuous creation of words such that they will evoke the desired meanings, at least in yourself. The "such that" part indicates the closure of the loop.

i think the closed loop runs in the other direction: we learn words in the context of perceptions--we know what the perceptions are and we learn the expressions to attach to them. when we want to convey similar perceptions, we choose among the expressions we know. there's very little creativity involved. (there's generally the **possibility** for creativity--but it doesn't happen v. often).

cheers.

--penni

Date: Mon Nov 02, 1992 3:54 pm PST
Subject: Misc replys

[From Rick Marken (921102.1300)] Bill Powers (921027.0500) --

Your post on "bandwidth" was a real gold nugget. There were some great "portable demos" for Gary Cziko's collection. I also think it reveals a potentially rich vein for research -- and research of the sort that can be easily done with "computer and mouse". I don't have time to start much right now but if there are any grad students out there who want to do PCT research, I would like to discuss some possible ideas for revealing levels of control via bandwidth studies (and modeling, of course).

Gary Cziko -- on making the "Blind men" paper more user friendly --

I agree with you completely! There -- how's that for a disturbance?

By the way, I got one review of the "Blind men" paper back from Psycholoquy -- the review was very nice and recommended publication in the electronic journal. But don't get excited; Harnad still wants to look at several other reviews before contaminating the e-ways with PCT. One thing the reviewer liked about my paper, by the way (N.B. Dr. Cziko) was that it was an attempt at integration of ideas in psychology! And he (she?) is RIGHT!

Bill Powers -- your post on gain in control was also a beaut. Thanks for taking the trouble to do this; even though I count myself as pretty PCT-wise, I always learn something new (or at least a new way of looking at what I already knew) from your posts.

A possibly new and inflammatory topic:

I was thinking of writing a paper called "Telling right from wrong"

The basic idea would be that PCT is a psychological theory (I would argue, the ONLY one) that can tell the difference between right and wrong. "Right" is just a real or imagined perception that matches a reference signal; "wrong" is a real or imagined perception that does not. Without references for perceptions, everything would just "be". But we consider some perceptions (experiences, situations, whatever you want to call them) better (more right) than others. This is a very basic experience for people; we are always evaluating the "rightness" of temperatures, pressures, speeds, directions, word pronunciations, word sequences, etc, etc. This is what PCT is about -- the fact that we can tell the difference between right and wrong -- at many levels of experience. PCT shows that the difference between right and wrong is determined by the setting of reference signals -- hence, by ourselves. Right and wrong are not "out there" -- they are in here; the current settings of all our references determine what is right. All controlled perceptions that do not match (pretty closely) these references are wrong.

So it is easy to test PCT -- just see if people think some things are right and others wrong. If they do, then they are control systems.

If you read the Bible, by the way, you will see that God is clearly a control system because s/he has definite (if peculiar) idea about what is right and what is wrong.

Best regards Rick

Date: Mon Nov 02, 1992 5:22 pm PST
Subject: Making right from wrong

[From Rick Marken (921102.1430)]

Addition to my last post:

Actually, I think I should call the paper "Making right from wrong" rather than "Telling right from wrong". There are theories that could be interpreted as explanations (actually, more like descriptions) of how people tell right from wrong (judgement theories, I guess, like signal detection theory, adaptation level theory, etc) -- what PCT does is add the observation that people act in order to make things right -- ie. they get perceptions to match their references -- and it explains how they do it (negative feedback control).

Best regards Rick

Date: Mon Nov 02, 1992 3:54 pm PST
Subject: Words, meanings, computations

[From Bill Powers (921102.1230)] Ray Allis (921102.0830) --

I moseyed downstairs just now with a vague notion of putting a question to speakers and listeners on the net: when I talk about meanings being perceptions, am I just talking to myself? And there was your post waiting, showing that at least one person gets the point perfectly, and in fact already understood it. Thank you, Ray.

I've been reading, or trying to read between bouts of frustration, Gerald M. Edelman's Bright Air, Brilliant Fire. It isn't Edelman himself who is the object of my frustration, but the pretentious hogwash (I sound like Tom Bourbon) that he picks apart so carefully. He refers to "The blend of psychology, computer science, linguistics, and philosophy known as cognitive science..." Everything he finds wrong with this occupation can be traced back to not knowing the difference between a word and the perception to which it refers. There is also a tremendous muddle about what's outside the brain and what's inside it. On top of that there is the whole "intentionality" fiasco, which wouldn't even exist if philosophers of science hadn't decided once and for all that purposive behavior couldn't possibly exist. Mary started reading this book first, and she kept exclaiming that Edelman really needs control theory. Now that I've got into it I agree (Mary often leads the way into works like these -- I haven't the patience to get through them unaided).

I'm also trying to write a piece on PCT, and having trouble getting it organized. One thought that came up while reading Edelman, and that your post reinforced, is that what we mean in PCT by "computation" is completely different from what mainstream thinkers mean by it. The mainstream concept of computation is strictly symbol-oriented. You convert the problem into symbols, then manipulate the symbols to arrive at a conclusion, a recommendation, a description, and so on. This is the sort of thing we learn to do in algebra classes, and also when we learn to program digital computers. This is not at all where PCT started.

PCT grew up in the years from 1953 to 1960, while the digital revolution was still gathering speed. Bob Clark and I were running computer models back then, but they weren't digital computer models (that would have been like trying to run a real-time simulation

on one of today's \$5 pocket calculators). We used a Philbrick analogue computer. This may have more to do with the difficulties of getting PCT across than anything else.

In an analogue computer you don't assign symbols to variables and then do computations with the symbols. You make one variable physical quantity directly represent another. Just by BEING of a certain magnitude, a voltage can represent the magnitude of something else like a force, a speed, a position. An analogue computer setup is an array of wiring connecting physical components. The components make one signal depend on other according to continuous rules of computation such as integration, differentiation, multiplication, addition- subtraction, thresholds, and limits. Those computations are not done with numbers, but with the properties of physical things. A capacitor becomes an integrator when put in a feedback path; it becomes a differentiator when put in series with an input path. A logarithmic converter is made of a diode. A multiplier is made of a diode bridge. An adder or subtractor is an operational amplifier with multiple proportional inputs and proportional feedback. If you want to express the fact that velocity is the integral of acceleration, you provide a voltage that stands for acceleration and make it the input to an integrator. The output voltage of the integrator, properly scaled, varies as the velocity would vary. Run that into another integrator and you get position, again represented as a continuously-variable voltage.

When you put signals into an analog computing network, you get signal variations everywhere in the network, not just at the nominal outputs. All of these variations have meaning, because they stand in continuous relationship to the inputs, according to rules that are the properties of the intervening devices. There isn't a symbol anywhere in the network: just voltages and currents behaving through time. Of course the user of the computer assigns symbolic meanings to the various signals: this is a muscle tension, that is an acceleration, the other is a sum of forces. But the computer itself operates without any such interpretations.

With this orientation, it's not surprising that I should have started out in the beginning thinking of the nervous system as operating directly with signals, not with the names of signals or symbols for their states. It's not surprising that I recognized the difference between a perception and a word that indicates a perception. The perception is just a signal that varies as the inputs to some analogue computing device vary. It doesn't have to have a name. It isn't a member of a category. All signals are alike. Their meaning is in the way they are derived from other signals, or from sensors.

What we were doing in those early years is now called "connectionism." Connectionism is just analogue computing making a comeback.

Ultimately, of course, we had to recognize the existence of words and symbols, and categories, and other things that seem to be perceived as discrete packages. But the same orientation still seems appropriate. A word is an activity in the brain; therefore it must be represented as a signal. A word indicates a meaning when it gives rise, through some continuously-present device, to another signal in a pathway that normally handles direct perceptions. Obviously we can't express the nature of the devices that connect the signals in terms of the signals; that would be like trying to draw a diagram of a radio that's playing music by using musical notation. Words and symbols are the signals handled by these devices; they aren't the devices. In the products of the brain we can recognize relationships that suggest processes in the brain, and the processes are created by the computing components, not the signals. The signals and their relationships, which are more signals, are only indications of the processes. I've been trying to express the understanding behind these words for 40 years, and still can't do it right. It's always been obvious to me that the conventional sciences are trying to explain how behavior works by looking at the CONTENT of a brain, the signals that happen to be present in it

-- and particularly those that happen to be present in the theoretician or experimenter. The experimenter tries to figure out how one signal behaves when another changes, but fails to realize that it is the functional connections, not the signals, that matter. When we observe a relationship between two perceptions, the important thing is not what that relationship is, or what the perceptions are, but that we ARE PERCEIVING A RELATIONSHIP. That, not the particular signals, is evidence about how the brain works.

This is particularly true at the "cognitive" levels. All of our digital-computer models and the complex mathematics that go with them are evidence about how the brain works, in terms of what it is doing. In themselves they are of little importance. An incorrect deduction is just as good evidence as a correct one. What we should be asking is not what cognitive calculations we can carry out or what new computations they lead to, but WHAT WE ARE DOING WHEN WE CARRY THEM OUT. We can't understand the computer just by looking at more computations. When we let the particular computations occupy our attention, we're missing seeing the principles of computation themselves. It's like starting out to understand the computer and getting hung up on a story that it's printing out. We have plenty of printouts already. We have to start looking for something else, something that explains the very existence of the printouts, not their content.

Well, you see what a bit of encouragement does to me. I don't know that I've managed to say these things any better this time, but it feels good to talk about them.

Is there anybody else out there besides Ray who gets the point?

Best, Bill P.

Date: Mon Nov 02, 1992 4:31 pm PST

Subject: Dropping Babies

[from Gary Cziko 921102.2118 GMT}

Bill Powers (921025.0800), as part of a very informative response to two questions of mine, said:

>Another "slowness" commonly cited is reaction time. The reaction time
>people usually think of is 200 milliseconds, the reaction time of a
>saccade or of a motor response to a sudden visual stimulus. This is
>far longer than actual delays in kinesthetic control systems. The
>delay in the lowest spinal-cord control loops is 9 or 10 milliseconds,
>and in brainstem loops only about 50 milliseconds.

This got me thinking of something I believe that is mentioned in the Roberston and Powers Introduction to Modern Psychology which asked something along the lines of whether, if you were holding a baby and you brushed your hand against a very hot surface, you would yank your hand away the way you would if you were not holding the baby (thereby dropping the baby). I believe the question assumes that you would not drop the baby. But why not? If the lower spinal loop acts so fast, how could you use the higher-level perception of baby-holding to prevent the spinal loop from doing its stuff and yanking your hand away?

I'm sure you've all had the experience of yanking your hand away from a hot surface before even feeling the sensation of pain. Does anybody have stories to tell of dropping babies (or similar examples of higher-level perceptions controlling faster-acting lower ones)?

--Gary

Date: Mon Nov 02, 1992 7:11 pm PST
Subject: Promoting PCT

[From Dag Forssell (921102.1700)]

A disk with updated CSGINTRO.DOC will be postmarked tomorrow for Gary Cziko and Bill Silvert. Gary will put it on the net and Bill make it available on the server. Pat and Greg have successfully (I could use it) created an ASCII file, uud.scr, which Bill will also put on the server. This allows MCI mail (and Bitnet?) netters with DOS machines to get the binary files. Instructions in CSGINTRO.DOC. Also, the marken.doc file will be updated, complete with ASCII figures. Instructions on how to subscribe to CSG-L have been added. Download the CSGINTRO.DOC and pass it around to friends.

A few respondents from this summer's mailings are still evaluating. We learned a lot from presenting our program and got some very positive feedback along with challenges. It is time to renew promotion efforts. Here is version 12 of our letter, addressed to a CEO. It should be more clear, better focused and shorter (cut from 4 to 2 1/2 pages). Comments please.

William Treval Powers
Independent Investigator #1
CSG Revolutionaries, Inc.
73 Ridge Road CR 510
Durango, Co 81301-8136

November 2, 1992

Dear Mr. Powers:

I am writing to introduce you to "Perceptual Control Theory" (PCT).

PCT offers a new explanation of how the mind works; why people do what they do. When you and your associates understand PCT and decide that it makes sense, you will look at yourself and others in a different way. You will know what can and cannot be done in harmony with others. This understanding and our methodology leads to cooperation. Productivity, quality, relationships and satisfaction will all improve.

As you know, a manager's function is to control and direct associates so tasks get done. The company performs, makes money and stays in business.

But we all dislike being controlled. We are self-directed and resent coercive authority based on position. A participant in our seminar asked:

You have shown us how to resolve conflict where I am a manager.
Can you show us what we can do peer to peer, where nobody has authority? (We did).

The functional requirement to *control* people who actually *control themselves* creates a management paradox. How it is resolved determines the level of cooperation and productivity in a company. This paradox can be resolved through a detailed study of how people direct themselves as autonomous "living control systems."

Working with others, each of us form ideas of how human beings "work." In the absence of a clear, valid and testable theory of psychology (many psychologists say that their theories and practices have nothing to do with each other), these ideas will vary widely in validity and effectiveness.

At the core of the design of your organization lie many such ideas about people. If you question and then change these based on an understanding of PCT, the impact on the design and function of the organization you are responsible for may be significant.

The new understanding of self-direction is based on clear, detailed and tested engineering concepts. Theory and practice go hand in hand. William T. Powers, who has developed the theory, writes in Living Control Systems, Vol II:

Over, please...

William Treval Powers

November 2, 1992

Page 2

Perceptual Control Theory explains how organisms control what happens to them. This means all organisms from the amoeba to Humankind. It explains why one organism can't control another without physical violence. It explains why people deprived of any major part of their ability to control soon become dysfunctional, lose interest in life, pine away and die. It explains why it is so hard for groups of people to work together even on something they all agree is important. It explains what a goal is, how goals relate to behavior, how behavior affects perceptions and how perceptions define the reality in which we live and move and have our being. Perceptual Control Theory is the first scientific theory that can handle all these phenomena within a single testable concept of how living systems work.

This 20th century understanding of Perceptual Control Theory (people always control their perceptions) provides a fundamental new insight that puts the problems that result from human interactions in a new light. It still looks the same, but the science of psychology can start over from a new concept and bring far better results.

Understanding people no longer has to be complex and confusing. PCT can be taught in simple form with a comprehensive management application in one day and in more detail with leadership applications in three.

The new perspective does not invalidate any wise common sense observation or practice. It does provide an enhanced understanding of seemingly intractable problems. It gives you new diagnostic tools and shows why cookbook rules for behavior (programs which tell you what to do under certain circumstances) do not always work.

Perceptual control is as incomprehensible and invisible at first glance to a person trained in cause-effect thinking (which we all are in our culture) as the idea that the earth revolves around the sun was to a person trained in the details of an earth-centered universe. A demonstration shows this clearly. Of course, this is only because it has never been noticed or explained. When you understand the principles, you will be aware of perceptual control in action in yourself and others. You will also notice that an understanding of PCT explains the various suggestions of cause and effect in people, just

like the understanding that the earth spins on its axis explains the illusion of the sun circling the earth.

I would like to describe this perspective so you get the point immediately. This is a difficult Catch-22 challenge, because this is a different concept altogether from what predominates in our world today. Until you have learned enough details, you cannot understand at all. I need a few hours in class to explain the functional elements of control and their interactions in simple step by step illustrations. 1

We provide step-by-step explanations of all aspects of Perceptual Control Theory with complete, clear illustrations and many participative demonstrations.

Participants learn a universal, simple and easy to remember methodology we call Teaching Effectiveness, based on the theory in an obvious way, that helps them think and deal effectively with real people problems as they occur.

Continued...

1 The enclosed report: _Control: What it is; where it applies_, provides a summary overview.

William Treval Powers

November 2, 1992

Page 3

Applications that follow directly from this approach are: conflict resolution; performance reviews; team development and non-manipulative selling. Role play practice and class discussion of experiences are provided for each of these.

Applications where the theory provides insight are: leadership understanding; goal structuring; vision & mission statements; a framework for total quality management.

A participant commented:

We have had quality management training from Philip Crosby. We were taught to focus on the problem rather than the person. You have plugged a big hole for me by showing me how to deal with people problems rather than avoid them.

People trained in the technical sciences appreciate the scientific approach and elegant simplicity of the theory, and everyone is able to begin applying the principles as soon as they understand the model and have had some instruction and practice.

I will be pleased to send you your personal selection of introductory material. You may request an audio tape with script to explain the nature, background and content of our programs. It includes an active do-it-yourself demonstration of the basic concept. The demonstration allows you to find out if your associates can recognize simple control in action. (I bet they can't). You may also request a report on results in industry, choose from a selection of papers and a very readable introductory book called: _Freedom from Stress_. The book is available on audio cassettes. This material will allow you to evaluate the validity of this information and its potential impact on your company.

You may find the demonstration enlightening and entertaining. This win/win program will increase the understanding and effectiveness of anyone who deals with people.

Sincerely,

Dag Forssell

Best, Dag

Date: Mon Nov 02, 1992 7:20 pm PST
Subject: RE: Words, meanings, computations

From Tom Bourbon [921102 -- 20:23 CST] - [Bill Powers 921102 15:11]

Bill asked if anyone else "got it" concerning perceptions, signals and language. Got it! And Bill, you absolutely should do the paper explaining the differences between "computation" in traditional AI and cognitive science, and analogue computing, which you, Clark and McFarland did in the beginning and which we try to emulate on today's digital PCs. That would be a great "bridge" article of the REAL KIND -- at least it would be for anyone on the "other side" who read it with the blinders off. (And if you want to sound like me when you comment on some of the over blown pablum that is out there, go right ahead!)

Regards, Tom Bourbon

Date: Mon Nov 02, 1992 7:38 pm PST
Subject: RE dropping babies

From Bill Powers (921102.2000)] Gary Cziko (921102.2118) --

RE: dropping the baby.

This story is like an urban myth: everybody knows about it but it never happened. Think in a little detail about how the burn reflex works. It moves the limb roughly away from the burned area. If you were holding a baby in your arms and brushed against a very hot object, which surface of the skin on your arm or hand would be burned? In what direction would the reflex carry the arm or hand? Exactly where on your arms would you have to be burned in order that the reflex would cause your arms to separate from the baby, thus releasing it so it would fall? How would this reflex cause your hands to relax their grip? What would you be doing with the unburned arm and hand?

It's easy to make up what-if challenges to control theory. But I think it's incumbent on the challenger to make the challenge real.

Best, Bill P.

Date: Mon Nov 02, 1992 7:49 pm PST
Subject: Creativity in Language Learning

[from Gary Cziko 921103.0240 GMT] penni sibun 921102.1200

I've been successful so far in keeping out of the language discussions, but this statment by Penni was just too hard to ignore:

>i think the closed loop runs in the other direction: we learn words
>in the context of perceptions--we know what the perceptions are and we
>learn the expressions to attach to them. when we want to convey
>similar perceptions, we choose among the expressions we know. there's

>very little creativity involved. (there's generally the *possibility*
>for creativity--but it doesn't happen v. often).

But how does the child learning language know which of the many available perceptions are being spoken about? The mother points and says "The cat is on the mat." The young child has to figure out which of her perceptions (which of the many that are potentially available) mom is referring to. Is she talking about her outstretched finger being itchy? Does she mean that the cat is under the piano? That the mat is under the cat? That the cat looks hungry? That the cat is tired? That I should play with the cat? That I should get something for the cat to eat? That there is dust under the piano?

Seems to me that there must be enormous amount of guesswork involved here, no matter what the linguists say. And LOTS of creativity on the part of the language learner who is continually inventing new solutions and selecting those that make most sense in his attempt to find consistent and useful relations between his nonlinguistic and linguistic perceptions.

--Gary

Date: Mon Nov 02, 1992 10:05 pm PST
Subject: Re: Creativity in Language Learning

(ps 921102.2200) [from Gary Cziko 921103.0240 GMT]

penni sibun 921102.1200

But how does the child learning language know which of the many available perceptions are being spoken about?

i don't know; do you? is this a question that only pct can answer?

Seems to me that there must be enormous amount of guesswork involved here, no matter what the linguists say. And LOTS of creativity on the part of the language learner who is continually inventing new solutions and selecting those that make most sense in his attempt to find consistent and useful relations between his nonlinguistic and linguistic perceptions. --Gary

sure. even as adults, we have plenty of experiences of having no idea what the heck the other person is talking about. the methods by which we (attempt to) figure it out are various: look at what they're (apparently) talking about; seeing what they say next; seeing what they do next; seeing how they respond to what we say or do next; ignore it and hope it doesn't matter; etc.. one method we can use is to think, ``i would say `blah blah blah' in what i think their situation is; maybe that's it.'' but i suggest that most of the time we're not generating language in order to understand language. in other words (so to speak), the ``new solutions'' created aren't new language, but new ways of associating our various ling/non-ling bits of knowledge together.

cheers. --penni

Date: Tue Nov 03, 1992 3:49 am PST
Subject: Opportunities for collaboration with other modelers

From Greg Williams (921103) Tom Bourbon [921102 -- 20:23 CST]

[commenting on Bill Powers 921102 15:11]

>And Bill, you absolutely should do the paper
>explaining the differences between "computation" in traditional AI and
>cognitive science, and analogue computing, which you, Clark and McFarland
>did in the beginning and which we try to emulate on today's digital PCs.
>That would be a great "bridge" article of the REAL KIND -- at least it
>would be for anyone on the "other side" who read it with the blinders off.

Sounds like a great idea to me, too. And Bill probably could find a number of possible co-authors for such a paper among the non-PCT researchers who publish in BIOLOGICAL CYBERNETICS -- many of these folks do modeling and simulations at exactly the same levels (NOT "traditional AI and cognitive science") as do PCTers, but without benefit of PCT/HPCT ideas. I'm saying that PCTers are only PART of the group on one side of the symbolic-only/non-symbolic-and-symbolic chasm, and that PCTers will find able and probably willing allies among non-PCT non-symbolic modelers in the polemic against a purportedly comprehensive traditional AI/cognitive science.

Best, Greg

P.S. Latest CLOSED LOOP ms. sent to Ed Ford yesterday. Look for "Conflict, Belief, Standards: Part I" in your hardcopy mailboxes soon.

Date: Tue Nov 03, 1992 8:34 am PST
Subject: Re: Right and wrong

From: Tom Bourbon [921103 -- 8:05 CST] -----

Re: "Right and Wrong" Rick Marken [921102 -- some time or other]

"Oh, ____! You have the ____est ____ (s) I ever ____ (ed)!" Followed either by, "I wish your ____ (s) could be a little more/a lot more/a little less/a lot less ____! Let me see if I can make your ____ (s) any ____er.", or by, "Your ____ (s) is/are just right!"

That little vignette was something I put together when I first read Bill Powers' BCP. I actually used it in some of my classes back in the 70s and early 80s. (No wonder my former colleagues were outraged when I began to teach about CST-PCT!) Right and wrong -- or "too __, too __, just right; and what am I going to do about it?"

That became part of a standard classroom gambit: Goldilocks vs King Midas. King Midas, of the Golden touch -- unidirectional causality with a vengeance; Goldilocks -- try it and if it isn't "just right," change it. But what isn't right for Goldilock is "just right" for either Mamma or Papa Bear. King Midas theories -- SR and IO/cognitive/neuroscience; the Goldilocks theory-- PCT. (That really tore it with my colleagues -- King Midas and Goldilocks?!) First-year undergraduates could pick up on that -- so could graduate students. Long after the first introduction of the comparison, students would spontaneously tell me about examples of Midas-think in theories from their other classes.

The Goldilocks gambit works. The "right or wrong" gambit will work. Maybe we could work on this together, or at least I could "consult." I still like Goldilocks vs King Midas

-- two falls out of three. Add in the three bears and it could be a tag-team match. (For those not familiar with televised "sports" in the USA, I am alluding to the theatrics of professional "wrestling.")

!Basta! Tom Bourbon

Date: Tue Nov 03, 1992 10:52 am PST
Subject: Re: Right and wrong

[From Rick Marken (921103.0930)] Tom Bourbon [921103 -- 8:05 CST] --

The King Midas/ Goldilocks comparison is INSPIRED. YOU should write it up -- if you haven't already. I get it completely. Midas is a PERFECT example of the S-R or output generation approach to life -- and is a good illustration of its consequences (if people really did it -- and, lord, do they try). Goldilocks is control in action -- she varies her actions to produce the RIGHT (for her) result.

>First-year undergraduates could pick up on that -- so could graduate
>students. Long after the first introduction of the comparison, students
>would spontaneously tell me about examples of Midas-think in theories
>from their other classes.

That's beautiful.

Then, if they are willing to learn a little algebra, they can get "The three blind men and the elephant". After they learn some modelling they could understand your two person interaction experiments -- which could be cast as "Androcles and the Lion".

Maybe we could get together and write "The child's garden of PCT".

I REALLY do love myths and fairy tales (that's why I read the Bible) -- but your use of them here, Tom, is absolutely sublime.

Best regards Rick

Date: Tue Nov 03, 1992 12:17 pm PST
Subject: Right and wrong linguistics

(penni sibun 921103.1100)

From: Tom Bourbon [921103 -- 8:05 CST] -----

Re: "Right and Wrong"
Rick Marken [921102 -- some time or other]

"Oh, ____! You have the ____est ____ (s) I ever ____ (ed)!"
Followed either by, "I wish your ____ (s) could be a little
more/a lot more/a little less/a lot less ____!
Let me see if I can make your ____ (s) any ____er.",
or by, "Your ____ (s) is/are just right!"

this is a great example of how language is often used. note the lack of linguistic creativity.

cheers. --penni

Date: Tue Nov 03, 1992 12:21 pm PST
Subject: Re: Right and wrong

From Tom Bourbon (921103 ---13:01) Rick Marken (921103 -- 11:44)

"A children's garden of PCT!" Let's do it!

Who knows, if we REALLY did it for children, we could subvert everyone and go directly for the next generations. They would simply laugh everything else off the stage.

Aesop's PCT Fables?

Regards, Tom Bourbon

Date: Tue Nov 03, 1992 1:08 pm PST
Subject: Re: right and wrong linguistics; Edelman

From Tom Bourbon (921103 -- 13:31 CST)

Subject: Right and wrong linguistics

>(penni sibun 921103.1100)

>> "Oh, ____! You have the ____est ____ (s) I ever ____ (ed)!"
>> Followed either by, "I wish your ____ (s) could be a little
>> more/a lot more/a little less/a lot less ____! Let me
>> see if I can make your ____ (s) any ____er.",
>> or by, "Your ____ (s) is/are just right!"

>this is a great example of how language is often used. note the
>lack of linguistic creativity.

And THIS (immediately above) is a great example of not thinking about the many different ways (in my naivety I would call them creative ways) students filled in those blanks. They took boilerplate and turned it into gold. Hmm, maybe Midas Models are best after all.

Could you give me a brief definition of "linguistic creativity?" Perhaps I missed it earlier. I think we are talking about different beasts.

Re: Edelman

Bill Powers (921102 -- late)

Thanks for the comments about Edelman's book, *Bright Air, Brilliant Fire: On the Matter of the Mind*. Andy Papanicolaou and I (and our Brazilian colleague, Luis Basile) have been wondering what it was about. We first saw it listed in an ad for the Natural Science Book Club, on the back of Science News, 12 Sept. 1992. Also in the ad was *Wet Mind* by Kosslyn and Koenig.

After reading the ad, I suggested to Andy and Luis that we should hurry to write a book on cognitive science -- *Dirty Minds*. That way the world could see how far "cognitive" science has come since the early Greek thinkers: On the market at the same time there would be three books that characterize the mind as Air, Fire, Water and Earth. Ah, the wonders of life in the decade of the brain!

Best wishes, Tom Bourbon

Date: Tue Nov 03, 1992 2:40 pm PST
Subject: Re: right and wrong linguistics

(ps 921103.1300)

>this is a great example of how language is often used. note the
>lack of linguistic creativity.

And THIS (immediately above) is a great example of not thinking about the many different ways (in my naivety I would call them creative ways) students filled in those blanks.

thanks for the ad hominem.

Could you give me a brief definition of "linguistic creativity?" Perhaps I missed it earlier. I think we are talking about different beasts.

i don't think anybody's defined it. in the thread that bill and i have largely been collaborating on, an issue was whether in language using/learning part of the process involves a large measure of creating afresh novel utterances. bill (i believe) was arguing for a great deal of creativity; i was arguing for minimal creativity. i do indeed think that most language production involved taking boilerplate and filling in the blanks. i also believe that restrictions on what can go in those blanks are few and elastic. however, insofar as the blank-fillers have their own structure, this structure is likely shaped by similar boilerplate.

cheers. --penni

Date: Tue Nov 03, 1992 3:47 pm PST
Subject: Re: right and wrong

From Tom Bourbon (ps 921103.1300)

Earlier you wrote:

>>>this is a great example of how language is often used. note the
>>>lack of linguistic creativity.

And I replied:

>>And THIS (immediately above) is a great example of not thinking
>>about the many different ways (in my naivety I would call them creative

>>ways) students filled in those blanks.

And you replied:

>thanks for the ad hominem.

No ad hominem intended -- I should learn to use some of the little tricks like :-).
All I meant was that I am not an authority on creativity, but I thought the ways students used the boilerplate were often very creative and that perhaps you hadn't thought of that possibility.

I had written:

Could you give me a brief definition of "linguistic creativity?"
Perhaps I missed it earlier. I think we are talking about different beasts.

You replied:

>i don't think anybody's defined it. in the thread that bill and
>i have largely been collaborating on, an issue was whether in language
>using/learning part of the process involves a large measure of creating
>afresh novel utterances. bill (i believe) was arguing for a great
>deal of creativity; i was arguing for minimal creativity. i do indeed
>think that most language production involved taking boilerplate and
>filling in the blanks. i also believe that restrictions on what
>can go in those blanks are few and elastic. however, insofar as
>the blank-fillers have their own structure, this structure is likely
>shaped by similar boilerplate.

Perhaps we ARE talking about different beasts. Do you mean that the term "creative" applies only if a person utters something totally new under the sun? You have confirmed that there was no earlier post of a definition for "linguistic creativity." That fact might account for some of the differences between the positions Bill and you have taken on the recent creativity-in-linguistics thread. Can you elaborate a little on your ideas in the previous paragraph? Perhaps what you say will help draw out the issues even more clearly.

Best wishes, Tom Bourbon

Date: Tue Nov 03, 1992 3:52 pm PST
Subject: Re: description

[From Rick Marken (921103.1500)]

penni sibun (921102.1200)


>> [From Bill Powers (921102.0830)]

>> You're getting close to the way I think about the link between words
>> and meanings (where, as you know, I define "meaning" as nonverbal
>> perceptions). The key word here is describe. What is a description
>> but a string of words that indicates some nonverbal happening?

>well, i think you might enjoy my dissertation, which argued just these

>things (though in comp. sci. jargon rather than pct jargon). you and
>i disagree on the relationship bet. words and meanings, but i think
>that's actually an open question. what is more important is that we
>focus on the process.

It seems to me that it is the process that is being disagreed on -- not the relationship between words and meanings. I think we all agree that words are just perceptions that stand for other (usually non-verbal) perceptions (although the word "verbal" does stand for a verbal perception). The question is -- what is the process by which the mostly non-verbal (maybe just call them "meaning") perceptions get created by a speaker or listener. In PCT the process is:



rml is the reference for the meaning perception that the speaker intends to produce. pml is the perceived meaning of the utterance. Any difference between rml and pml is error, e, that leads to variation in ws, the word strings being generated; this is the process that I'm sure Bill was referring to when he talked about "creating words"; the function transforming e into ws may be quite deterministic; the word generation process is not necessarily like pulling words out of a hat (though it could be). The words (ws) are part of a controlled variable (cv) that is also influenced by disturbances (which I call c for "context"; I imagine that one of the main disturbances to the variables controlled in speech is context like other things just said, inflections, other things going on, etc, etc.). The cv is continuously transformed into a "meaning" perception by the speaker (pml) and the listener (pm2).

The important part of the process in the PCT model of language is that it is closed loop -- the words that represent the meaning (ws) of what is being said are generated by an error signal -- not by the intended meaning itself, rml. In fact, there may be many different ws's that lead to the pml that matches the intended meaning; the appropriate value of ws that produces the intended meaning, rml, depends on disturbances (the context in which the words are said). So the PCT model of meaning says that what you are generating when you are producing meaningful utterances is not ws but pml -- you are controlling a meaning perception, not a word string output. The value of ws depends on c, not on rml (the intended meaning).

This model, is, of course, only a sketch. To make it work, you need to be able to convert cv into pml; not a trivial task, maybe? And you must be able to convert rml-pml into an error signal that preserves negative feed-back. You also have to know what cv is, of course, so you have to do a little testing for controlled variables. But that's what linguists are for, right?

Was this the process used in your model?

Best regards Rick

Date: Tue Nov 03, 1992 4:07 pm PST
Subject: Re: Language

(penni)

[From Bill Powers (921029.2200)]

Penni Sibun (921029) --

That injury must be maddening to a person who writes a lot. Is there anything that can be done to make it better? How about one of the new non-keyboards that you hold in your hand, pressing buttons in pairs? This would at least require different motions and positions. I read about it in Byte a few months ago; it seems to be learnable.

that would probably be good; i've heard of them. a workable voice interface w/b even better. hopefully when i get a real job, i can ask for better computer-using support (or afford it).

>language is not sentences in the same sense that space is not ether.

You've made your point that language doesn't necessarily come in sentences.

great. i would like to think so.

A lot of people speak in sentences, if they're not politicians. Is this just an affectation?

no, it's a misperception. you take a videotape of 1 or more people talking, try to transcribe everything they've said, and try to insert sentence boundaries. (nb, it's *cheating* to only insert *some* sentence boundaries. it doesn't prove anything if you can mark off, say 65% of the text as being included in some sentence or other. either language is made up of sentences or it isn't.) if you can manage that, then you could get a nobel prize for linguistics.

cheers. --penni

Date: Tue Nov 03, 1992 5:03 pm PST
Subject: Re: description

(penni sibun 921103.1600)

[From Rick Marken (921103.1500)]


The question is -- what is the process by which the mostly non-verbal (maybe just call them "meaning") perceptions get created by a speaker or listener.

no, the question was, at some point, how speakers relate ``meanings'' and ``words.'' i guess i'm almost as uncomfortable by the implication that ``meanings'' are constantly created as i am by the suggestion that ``words (language)'' are constantly created.

In PCT the process is:

Speaker

Listener



thanks for your diagram and explanation. i found them v. helpful.

This model, is, of course, only a sketch. To make it work, you need to be able to convert cv into pml; not a trivial task, maybe? And you must be able to convert rml-pml into an error signal that preserves negative feedback. You also have to know what cv is, of course, so you have to do a little testing for controlled variables. But that's what linguists are for, right?

right. i've spent a lot of time trying to characterize the stuff you call c and cv.

Was this the process used in your model?

well, *you* have my thesis and journal paper; you tell me. really--i'd be interested to know.

i think the major difficulty in mapping bet. yr model and mine is the ``intended meaning'' part. this is a basic sketch of my model: i cast the issue as deciding what to say next. there is an enormous and complex context that constrains the available choices. 2 of the most important elements of the context are

1) what linguistic structures are licensed at this point and 2) what is available to be talked about next; another way to characterize this might be 1) where are we in the linguistic structure (ws) and 2) where are we in the structure of ``meaning(s)'' (?) we are talking about. ignoring the rest of the context for the nonce, what is said next is determined both by 1) and by 2). this implies that it is *not* the case that the ``meaning'' is selected and then the linguistic structure is found to fit it, and the next bit of language is produced. there may be 4 things you could talk about next, but given what language you've produced so far, only one of them continues what you've said felicitously. it's that felicity that determines the ``intended meaning'' in this case; the ``intended meaning'' and the linguistic structure are determined *simultaneously*. how would you represent that in pct?

cheers. --penni

Date: Tue Nov 03, 1992 5:12 pm PST
 From: Control Systems Group Network
 Subject: sentences

[Avery Andrews 921104.1205) (Penni Sibun (921029))

>no, it's a misperception. you take a videotape of 1 or more people
 >talking, try to transcribe everything they've said, and try to insert
 >sentence boundaries. (nb, it's *cheating* to only insert *some*
 >sentence boundaries. it doesn't prove anything if you can mark off,

>say 65% of the text as being included in some sentence or other.
>either language is made up of sentences or it isn't.) if you can

This way of putting makes an important point, that sentences can play a role in language-use without ordinary language actually being made up of them in any concretely useful sense. More generally, the big problem in getting a useful PCT slant on linguistics is to isolate what's actually true and useful about current linguistics (I think there's quite a lot it is), & separate it out from the gunk (of which there is also quite a lot).

Avery.Andrews@anu.edu.au

Date: Tue Nov 03, 1992 7:10 pm PST
Subject: description

[Avery.Andrews (921104.1349)] (Rick Marken (921103.1500))

>The important part of the process in the PCT model of language is that
>it is closed loop -- the words that represent the meaning (ws) of what
>is being said are generated by an error signal -- not by the intended
>meaning itself, rml. ...

In a sense, something like this is already built into Salix. The entities and relations in the KB that is to be expressed have a bit which marks 'already said', and the generation process is organized so as to go on until everything is so-marked.

The problem is that the reference level (the whole KB) & the controlled perceptions (the portion of the KB marked 'already said' are not continuous-valued vectors (or something like that, I haven't gotten a real grip on the bandwidth issue), which are what is wanted in neurologically plausible modelling, but elements in a discrete lattice. This is a really big problem -- people just don't know how to make connectionist systems do the sorts of things that speakers do.

One possible way forward for PCT might be to try to get a grip on what more of the things are that are being controlled for in linguistic communication. Salix just says everything it knows, for example, but people try to be a bit more selective, and sometimes manage to be. Given that somebody knows various things, there's still a lot of problems in explaining how they come to say exactly what they do say. My guess is that this would be a more promising area than trying to produce a PCT theory of syntax or lexical choice.

Avery.Andrews@anu.edu.au

Date: Tue Nov 03, 1992 9:39 pm PST

i.n.kurtzer(921103.2315)

a interesting thing happening while i was studying today that i thought might be of interest to any pct-ers out there; i was concentrating on reading a book and reached for my cup of coffee. the cup was in front of the book and its bottom half hit the book because i wasn't paying attention to its course; instead of backing up the cup and raising it higher i put more effort into pulling it towards me, then "just in the nick of time" (i hate these cliche's)

a switched the cups direction and stopped it from spilling on me. i thought this example was relevant to a experiment bourbon showed me once where the person lost control of a cursor or the way his//her actions controlling the cursor changed and the person increased his/her output in that direction. positive feedback i suppose? i'm just glad that that increase in output of mine was checked negatively quick enough.

sorry about the grammatical mistake, but the computer doesn't allow me to

like this skip of two line and the singular form of mistake above correct by going up to what i have written ,#%^*!!

just your average steelworker
quality you can trust
i.n.kurtzer

Date: Tue Nov 03, 1992 10:04 pm PST
Subject: description, interaction

[Avery Andrews 921104.1654]
(Rick Marken (921103.1500))
(penni sibun (921102.1200))

Here's an aspect of the PCT-ification of Salix that I've been overlooking for a while. Salix as it stands simply marks information as said when it thinks it's said it, but what's really going on has to be a bit different. In the first place, for adult conversation, the things `to-be-said' should really be taken as things `to-be-known-by-the-hearer'; so that there reference `signal' is what the hearer is wanted to know, the perception what the hearer is `perceived' as knowing, where a person's being in the presence of X being said is normally taken as a sufficient basis for concluding that the person knows X (there must be a better way to say this, but it eludes me for the present).

The point of all this: suppose two people, say a mother and father, are recounting some incident (what their kid did) to a third person. They can do this without both of them just reeling off everything they know about the incident, since one person's presenting an aspect of the incident cancels that contributor to the error signal for both of them (I think that what is going on is that one person takes the role of primary narrator, the other filling in with such bits as the first leaves out). Having the loop closed throught the environment should make it easier to explain how cooperation works.

With children, I'm not so sure about all this. My impressions are that (a) they just want to say what they want to say (and definitely don't want anyone else to say it) (b) aren't really focussed on conveying new information to a hearer--savoring common knowledge seems to be more like what they're on about.

Avery.Andrews@anu.edu.au

Date: Wed Nov 04, 1992 12:50 am PST
Subject: origin of control

[Hans Blom, 921030]

Some time ago I reported to you my interest in the origin of control: what is the simplest control system possible. The link between physics and psychology continues to fascinate me. Let me give you a summary of some further thoughts and on some more simulations that I have done. There may be some important implications for organisms, even human ones, as well, but there I must speculate (i.e. reason by analogy).

The birth of control

Around 1827 the herbalist Brown looked through his microscope and discovered how pollen submerged in water made strange, ever changing random zigzag movements. Not before some fifty years later this random motion of micron-sized particles in water was understood in terms of the movement of water molecules produced by what was known as 'thermal energy'. Initially it was thought that the movements were caused by water molecules pushing against the particle, but that view was too simplistic: the particle is much too large and heavy for a single molecule to be able to move it. The movement is statistical in nature: molecules bump into the particle from all sides equally on average, but over a short time it cannot fail that the number of pushes from one side is slightly larger than that from the opposite side. The result is a random fluctuation in the number of bumps against the particle from all sides causing a particle to move in random-sized steps in random directions.

In 1905 Einstein himself gave a theoretical foundation for Brownian motion by relating the particle's thermo-kinetic energy to the drag forces that arise due to its movement in the fluid (as had already been described by Stokes). Einstein provided us with the formula

$$\overline{(dx)^2} = \frac{k.T}{3.\pi.\eta.r} . dt$$

where

$\overline{(dx)^2}$ = the (average) square of the particle's displacement during a time dt; the direction of the displacement is random
k = Boltzmann's constant
T = absolute temperature
pi = 3.14159...
eta = the fluid's viscosity
r = the particle's (average, effective) radius

Note that, due to the vector addition of random displacement vectors, the average size of each step (we sample at multiples of a fixed time interval dt) is proportional to the square root of dt. But much more important in this context is the relationship between the average step size and the particle's radius. Let us investigate what happens when the particle's radius can vary.

Assume that we add some chemical that causes the particle to swell to the water. It is not particularly difficult to find such a chemical for many types of particle; even a great many not too small organic molecules show a change in shape in the presence of a large variety of other, small or large, molecules. Everything else being the same as before, the now larger particle will move more slowly. We can also create a concentration gradient if we drop a slowly dissolving pellet or trickle a micro-flow of the chemical into the water. The particle will now be largest and thus move slowest in the area where the concentration is highest.

If we pick some mathematical formula to describe the relationship between the chemical's concentration and the particle's radius [$r = f(c)$], and another formula for the spatial distribution of the chemical's concentration [$c = g(x, y)$ in two, $c = g(x, y, z)$ in three, or $c = g(x, y, z, t)$ in four dimensions], we can model the particle's movement over time. What do we find?

To simplify matters, I did my simulations in two dimensions. I varied the functions f and g but unvaryingly found that if only the radius of the particle increased with the chemical's concentration, the particle persisted for a longer time in those regions where the concentration was highest. As was to be expected. Given an unchanging concentration profile $c = g(x, y)$, the persistence of course depends on the nature of the relation $r = f(c)$ between concentration and radius. The course is computed by choosing an initial position (x, y) and by adding, at each sampling time $i \cdot dt$, $i = 1 \dots N$, an increment of size

$$dx = \sqrt{\frac{k \cdot T}{3 \cdot \pi \cdot \eta \cdot f(g(x, y))}} \cdot dt$$

in a random direction. For a wide variety of choices for $f(c)$ and $g(x, y)$, we can distinguish between the following simulation results:

1. If the particle's potential radius increase is nil, we see normal Brownian motion. It is a characteristic of Brownian motion that dx increases without bound. In due time, the particle disappears from sight.
2. At small values of the particle's potential radius increase, the particle will have a tendency to move toward the maximum of the concentration, but the thermal energy overcomes this tendency and more or less quickly moves the particle out of sight.
3. With larger values, the particle will not disappear from sight if in its initial position it can 'sense' a strong enough concentration. Although its random movements sometimes may take it relatively far away from the concentration's maximum, the particle does not get lost. You could say that the position is 'controlled', but control is poor.
4. With even larger 'gain' values, the particle stabilizes its position increasingly better near the maximum of the concentration. We now start to have a 'good' control system, where Brownian motion is only a random disturbance that causes jitter in the particle's position. This seems to be the case, however, only with rather contrived functions f , which seem quite unrealistic for large molecules or other 'dead' particles like pollen. Maybe small bacteria that have extensible but not otherwise controlled cilia do navigate this way.

Conclusion at this point: control is not an all-or-none phenomenon; there is a plastic progression from no control at all to almost perfect control. In simple control systems it seems difficult to demarcate sensor(s), comparator(s) and actuator(s).

Hypothesis at this point: in all organisms we will encounter instances of this progression from no control at all to (almost) perfect control -- from being swept along by the winds of fate to almost perfect control over e.g. the chemical composition of body cells.

Speculation at this point: is control necessarily more perfect at

low levels such as the chemical composition of body cells than at high levels such as self-respect or happiness?

One feature that is easy to add to the simulation is a (non-zero) setpoint: let the particle's radius start to decrease again when the concentration gets too high. This will take care of positioning the particle at a location where the concentration results in a maximum size for the particle. If the particle were a primitive organism, it would position itself at its optimum feeding location, between hunger and poisoning.

Now do assume that the particle is an organism and that the laws of evolution are operative. In that case descendants of the organism will, all else being the same, be selected on the basis of their f -function. A less appropriate f and Brownian motion will take them away and let them, if not starve then at least grow and multiply more slowly than their better designed cousins. Evolutionary laws will, slowly by slowly, design ever more appropriate f 's and thus ever more effective control systems.

Note that Brownian motion should not be abolished by a primitive organism. It has a very important function: it is responsible for survival whenever food gets exhausted and new food must be found. An extremely hungry organism is presumably at its smallest and thus moves fast (see Einstein's formula above). It need not expend any energy at all in its search for food; the environment's thermal energy does the job. Effectively, the organism could 'hibernate' while water molecules move it around; this is in fact what happens in numerous viral and bacterial species. Brownian movement thus sets a 'random search' for more food in action. While of course this is the only search strategy these most primitive organisms can have due to their lack of sensors, actuators and memory, from an engineering perspective random search is an effective search strategy when no knowledge whatsoever about the environment is available.

Pure chance (Brownian motion) is thus extremely purposeful for the survival of these primitive organisms. Not from the point of view of the individual organism, of course: although it has a well-defined goal (get to the food!), there is no directed action that makes sense. No food source is in sight, no gradients or derivatives of any type are available. Calculation nor deliberation, intention nor willpower provide, literally, a course of action. But then it is as if physics mercifully provides a backup: "If you don't know what to do, I do. I'll even do it for you". Providence.

Is Brownian motion relevant for higher organisms as well? The answer is a definite no. Once the organism becomes too large, thermal energy cannot significantly move it around anymore. A random search mechanism might be equally advantageous, however, in situations where control does not apply because the objective is nowhere in sight and thus no course of action can be computed. Could it be that the antithesis of control is not standstill but random action? Could 'random search' be the supreme endeavor to break stalemates? Is that what trial-and-error learning (reorganization) is all about? Could it be that -- in addition to reorganization or 'passive' learning -- evolution has taken care of some forms of 'active' random investigation in higher organisms as well? Is panic a purposeful 'higher' alternative for Brownian motion? How about play?

Date: Wed Nov 04, 1992 5:05 am PST
Subject: plant nervous systems

There's a kind of survey article in New Scientist for 10/17/92 (29-33). It's based on The Action Plant (Blackwells, July 1992), the author is a television producer, and there do not appear to be explicit bibliographical references to the literature, but with the implied caveats it might be useful. I haven't yet had a chance to read it.

Bruce bn@bbn.com

Date: Wed Nov 04, 1992 12:35 pm PST
Subject: Control in nature

[From Dag Forssell (921104-1)] Re: Hans Blom, 921030 origin of control

What a magnificent posting. You challenge the imagination.

In regards to control and physics / chemistry, I recall mention by Bill over a year ago of chemical processes that stabilized in a manner suggestive of control. Could anyone provide insight into this. Any references? Any descriptions accessible to the lay person?

Thanks, Dag

Date: Wed Nov 04, 1992 2:22 pm PST
Subject: replies on connectionism, language

[Martin Taylor 921104 13:40]

I'm trying to catch up on the backlog, so here are a few comments on different postings, mostly from Bill Powers and mostly on language, over the last week or two. I don't expect to be away again for any significant period until Spring, so maybe I'll get around to some earlier postings that I have been saving for comment when time came available. (Statistics for you, Rick? Maybe soon.)

First, before getting into the language stuff:

>From Bill Powers (921102.1230)
> What we were doing in those early years is now called
> "connectionism." Connectionism is just analogue computing making a
> comeback.

I disagree strongly. In the past I have accused you of promulgating a symbolic version of HPCT as opposed to a connectionist view of HPCT. I still think this is what you do. You give each ECS a specific function, appreciably different from that of each other ECS, and describable in symbolic terms. A connectionist view of HPCT is what we are trying to achieve in the long term in the Little Baby project (though admittedly the present LB has far too few ECSs to qualify). Connectionism depends at root on the distribution of representation, not on whether some elements have a continuous rather than a discrete range of variation. Your version of HPCT isolates the responsibility for the control of particular percepts to particular ECSs. We would not do that, but would share the responsibility through overlapping quasi-modular groups of ECSs. Your version is to connectionist HPCT as classical AI is to connectionist classifiers. And it leads to related kinds of problem.

=====

>From Bill Powers (921027.0700)
>I want to make a modest push to make sure our linguists really truly
>get this point. They may actually get it, but I don't think we're
>hearing the result. Sentence construction is not construction of some
>object out there, or in some vague conceptual space; it's construction

>of an input, a perception. The mere fact that we know of a sentence
>shows that it's a perception. The same goes for grammar, for any
>regularity we PERCEIVE in language. The relationships between
>different levels of analysis of language are relationships among
>levels of perception, not levels of output production. We have to
>guess what the production processes are, because they're outputs and
>we don't perceive outputs. We perceive only their perceptual
>consequences: language is perception.

Yes, but in most cases of language production, I would have thought that most of the perception was imagined. There is very little overt correction of ongoing language production, though there may well be correction based on the talker's perception of the effect of the production on the hearer. The actual output utterance seems usually to be produced accurately (at the level of the phrases--I'm not talking about the ongoing control of articulatory perceptions).

What we are into here is a bit of a problem with the relation of imagination and planning. The lower-level control processes in speech are usually good enough that few if any real-world disturbances affect the production of intended words; it could almost be a straight outflow process! So an utterance could, in principle, be preplanned by imagining the effect it would be likely to have. The controlled perceptions would be at higher and lower levels, where the real-world disturbances are important.

>

>From Bill Powers (921028.0630)

>(Penni Sibun)

>>i would argue (and do, very strongly) that our perceptions of
>>sentences and grammars have next to nothing to do with how we use
>>language, either understanding it or producing it.

>

>I agree in the sense that we use words in any way needed to evoke
>meanings in ourselves that we hope are similarly evoked in others. But
>I disagree with your extreme position, because it's clear that people
>DO make well-formed sentences and perceive them as such; you did so in
>the sentence above. We have concerns not only with conveying meanings,
>but with observing grammatical rules and customs. Sometimes it is only
>the grammatical form that lets us figure out what someone else's
>communication should mean. Syntax and grammar are aids in narrowing
>down possible meanings; we control for them the most carefully when we
>are most concerned with getting an exact meaning across.

The Layered Protocol position on syntax is that it is a conventional way in which the functions of different elements of a message at one level can be determined at that level. A message that is properly formed syntactically is one in which sender and receiver of the message agree on the ways in which these functions are delineated, whether it be at a phonetic level, the level of phrases, or the level of arguments. In this sense, syntax is the one place in language where "encodingism" might apply.

It is assumed by the talker, for example, that a raised pitch, lengthened syllable nucleus, and increased amplitude will identify to the listener an element that demands special attention (it might be new information, or might contrast with something the talker believes the listener to believe). These are conventionalized cues (I believe that in Danish pitch is lowered to have the same effect). At a much higher level, one might begin an argument by saying "You know that A" rather than simply "A," which might

indicate that the talker was asserting A rather than using it as a prerequisite for the assertion to follow.

I hesitate to say that syntax has any meaning in itself (a position not taken in the quoted chunk, but mentioned in some postings in this thread). Syntax is any form of redundancy (in the Shannon sense) that is incorporated in the structure of a one-way message. It helps assure the correct transformation of messages from one level to another.

>

>From Bill Powers (921029.2200)

> I agree, language theory has to handle spoken language. It may be
> that written language is somewhat artificial -- but if you look at
> transcripts of spoken language, you have to wonder how much
> imagination the listeners have to use to get a meaning out of it,
> much less the intended meaning.

I don't think written language is in the least artificial. The mistake, to me, is in identifying written language as a way of making a permanent record of what might be spoken. One can use writing that way, more or less, if you don't mind missing a lot of what is spoken, but written language is really an independent system, and it may well have originated quite apart from the idea of recording speech, in agricultural and calendric records.

Written language is different from spoken language in various ways. The most important difference is in the timing of the feedback that lets the talker or writer discover the effect of the speech/writing on the recipient. The talker can modify the effect as perceived on a time-scale of seconds, whereas the writer (even in e-mail) can do it only on a time scale of hours to months. This means that the written language must be structured so that its internal patterning enables the message to resist many kinds of difficulties. Written language must have a strong syntax, or be so clearly situated in context that its intent can be derived without the syntax. It must have many synchronizing tricks that enable the reader to take parts in relative isolation and interpret them even if preceding parts have been missed. Spaces between words, paragraph boundaries, conventions such as "a sentence must have a subject and a verb", cues to the argumentative function of a sentence, all are synchronizing syntactic tricks that have little or no analogue in conversational speech.

>from Gary Cziko 921103.0240

>But how does the child learning language know which of the many available
>perceptions are being spoken about? The mother points and says "The cat is
>on the mat." The young child has to figure out which of her perceptions
>(which of the many that are potentially available) mom is referring to. Is
>she talking about her outstretched finger being itchy? Does she mean that
>the cat is under the piano? That the mat is under the cat? That the cat
>looks hungry? That the cat is tired? That I should play with the cat?
>That I should get something for the cat to eat? That there is dust under
>the piano?

Mothers don't start off with complex sentences like that. "Motherese" is a label for the way they tend to talk to babies. Mother might say "Cat, cat" (more probably "Pussy, pussy") while letting baby stroke the cat. Baby says "Ca'", mother says "yes, cat" and smiles.... Baby learns that this happens only when pussy is in view. And so forth. Two-word phrases that link what Avery calls a thing and a situation (if I remember rightly) come next. It takes a while for complexes like "the cat" and "the mat" to be

linked into a single situation. By that time, the intersection of many different experiences with cat->"cat" and mat->"mat" have probably rubbed off many of the rough edges of the relation, but it is still not the same relation as adults take for granted. Children do a lot of what is called "overextension" and "underextension." The word "Clock" may be used to apply to a wall clock and a plate and a round hat, but not to a tall clock (a grandfather clock, in some dialects).

It is not only young children that must "figure out which of her perceptions >(which of the many that are potentially available) mom is referring to." We all have to do it when confronted with a new word for some concept, and we may never be sure that we have the same mapping as anyone else does. Does "democracy" involve a multiparty system of reconciling antagonistic sets of views? Could a two-party system ever be considered "democratic?" Some people would say "yes" to one question and "no" to the other, some would say "yes" to both, some would say "no" to both, but all would use the word "democracy" quite happily in a discussion, possibly without ever discovering that the parties had quite different answers to these basic questions about the "meaning" of the word.

=====

On "meaning" generally, a topic which has come up quite a bit in the last couple of weeks.

"Meaning" is as meaning does. It is pointless to argue about the meaning of "meaning", because it has many. Words out of context can often be identified as labels for concepts of one kind or another, but these definitions can only be guides to assist in a conventionalized use of the words. It may not be too much exaggeration to say that one can invent scenarios in which any word can take on any predetermined meaning (i.e. have any predetermined effect) at all.

In a trivial way, two people can agree that if one says "toast," the other will throw a football through the living room window. Does "toast" then label this bizarre piece of destruction? In a way, it does. Less trivially, without previous agreement, one person might manoeuvre the other (I hesitate to use the word "control") into doing just that when "toast" is mentioned (hard, but possible). Does "toast" then mean "throwing a football through the window?"

In Layered Protocol theory, there is no concept of "meaning." All there is is the effect the message at one layer has on the effect of a message at another layer, and eventually the effect that the transaction has on the perception the talker has of the listener's beliefs and vice-versa. All there is is the control by each partner of their perception of the other's beliefs. "Meaning" is the change wrought by the transaction--the error-correcting actions that bring the perception of the other's beliefs to match each one's reference levels for those beliefs.

In this sense, then, "meaning" can no more be an absolute than action can be behaviour. Each (behaviour and meaning) is the control of perception, but "meaning" is attached to the behaviour we call language.

Rick and Tom (and others?)

The idea of PCT teaching fables is fabulous. PCT is obviously hard to get across in words, like Sufi. But, like Sufi, it may be possible to approach obliquely through fables that show how non-PCT approaches dissipate into thin air when blown away by the

truth of PCT. The Midas touch of Goldilocks would be a good starter! Where is the PCT Nasruddin?

Martin

Date: Wed Nov 04, 1992 4:41 pm PST
Subject: language, origin of control

[From Rick Marken (921104.1400)]

It's over. The 12 year nightmare of greed, division, deceit, and selfishness is kaput (for the time being). I think we finally have an 11 level system in there -- instead of the usual 10 level (at best) variety (they walk and talk OK but don't know why). An 11 level system understands that ALL people are control systems -- not just rich, protestant, heterosexual, white males. Maybe we can try cooperation (community) for a change; I hope, I hope.

There. Got that off the ol' chest.

penni sibun (921103.1600) on the relation between her model and pct model

>well, *you* have my thesis and journal paper; you tell me. really--i'd
>be interested to know.

I hate to say this, but I have a lousy filing system -- so I don't know where they are. But I do know that whatever you did it struck me as being interesting enough to merit my inviting you onto csg-1 -- and I'm sure glad you joined.

>i think the major difficulty in mapping bet. yr model and mine is the
>``intended meaning'' part.

I agree -- intentions (reference signals) ARE the ONE BIG difference between PCT and ALL other models of living systems.

> this is a basic sketch of my model: i cast the issue as deciding what to say
> next.

Doesn't deciding mean comparison of one thing to another? In PCT there is a continuous comparison of intended to actual perceptual result.

> another way to characterize this might be 1) where are we in the
>linguistic structure (ws) and 2) where are we in the structure of
>``meaning(s)'' (?) we are talking about.

In order to be able to answer these two questions, doesn't the model have to be able to perceive the state of ws and "where we are in the structure of meanings"? That is, 1) and 2) describe perceptual variables, right?

> ignoring the rest of the
>context for the nonce, what is said next is determined both by 1) and
>by 2). this implies that it is *not* the case that the ``meaning'' is
>selected and then the linguistic structure is found to fit it, and the
>next bit of language is produced.

Your second sentence is not a correct description of the PCT model of meaning generation as I described it in my diagram; the PCT model says that a linguistic structure (ws) is generated which moves perceived meaning into a match with reference (selected or intended) meaning. You say that 1) and 2) [the questions quoted above] determine what is said next. To me, it sounds like the perception of 1) where you are in the linguistic structure (call this p.ws) and 2) where you are in the structure of meanings (call this p.m) determines what you say next (call this o). So, your model seems to be $o = f(p.ws, p.m)$ -- definitely not a control model. But you may have a control model without knowing it; what you say next may be ws and m rather than o -- so $(p, m) = f(p.ws, p.m)$ AND $(p.ws, p.m) = g(p, m)$ -- closed loop.

> there may be 4 things you could
> talk about next, but given what language you've produced so far, only
> one of them continues what you've said felicitously. it's that
> felicity that determines the ``intended meaning'' in this case; the
> ``intended meaning'' and the linguistic structure are determined
> *simultaneously*. how would you represent that in pct?

I think the only way you can determine whether the one in four things you could have said had a felicitous result is if you have some specification in the speaker of what constitutes such a result; in PCT this specification is the reference signal (r.m) in my diagram). If the result of the 1 of four outputs is a (perceived) result that matches the reference, then that output produces a felicitous result; felicity, in PCT, is the nearly complete absence of error (difference between reference and perception). In PCT, the linguistic structure (output) and intended meaning (a perception of meaning that corresponds to the reference) do not occur "simultaneously" but nearly so; anyway, these variables are varying simultaneously, with a phase lag.

Hans Blom (921030)

> Conclusion at this point: control is not an all-or-none phenomenon; there is a plastic progression from no control at all to almost perfect control.

Nobody ever imagined that control was all-or-none. The gain of a control system can go all the way from - infinity (perfect control) to 0 (no control) to + infinity (perfect amplification). If, however, the LOOP GAIN of a control system (the algebraic product of amplification factors around the control loop) is ≥ 0 then the system does not control; there is NO control AND there is no control system. If the LOOP GAIN is < 0 then you've got a control system (though a piss poor one if the gain of the system is not large). So there is a QUANTUM difference between systems that can control AT ALL and systems that just do not control; the difference is that between a closed loop (with negative loop gain) and an open loop (with 0 loop gain) or closed loop (with positive loop gain). Only a closed negative feedback loop has even a chance of exhibiting control.

> Hypothesis at this point: in all organisms we will encounter instances of this progression from no control at all to (almost) perfect control -- from being swept along by the winds of fate to almost perfect control over e.g. the chemical composition of body cells.

This is ambiguous. If you mean there is a progression from no control to control within a single organism then the hypothesis is known to be true; when I learn to control a cursor on the computer screen I go from virtually no control (disturbances are completely effective) to perfect control (disturbances have no effect). If you mean there is a progression across organisms then I think it is false; there is no organism that has

(almost) no control -- except of variables that it cannot perceive. If a system is alive, it controls (something) and it does it well. The transition from a non-controlling to a controlling system is quantal -- the system is either closed-loop with respect to some variable(s) or it ain't. If it is closed loop, then it controls that variable if the loop gain is $\ll 0$. The transition (in evolution) from non-controlling to controlling systems must have been sudden -- at least the process of going from an open to a closed loop structure and then (if necessary) going from a positive to a negative loop gain.

Best regards Rick

Date: Wed Nov 04, 1992 5:12 pm PST
Subject: Re: Squirrels, airconditioners

[From Jeff Hunter (921104)] Re Bill Powers (921028.0630)

Thanks for your comments on my squirrel posting. I think I'll try my point again with exposition first, and examples to follow.

Your contention is that "proper" control systems have loop gains of at least -5 (preferably -10 to -100). I believe that this rule-of-thumb of yours may be a result of the specs of the systems you designed. I also believe that this is too constrained a definition to describe the majority of control systems in living organisms.

A "good" control system keeps its error within a specified distance of zero for a specified percentage of the time for a given class of disturbances.

The control systems you describe generally have a small tolerance for error, quick response time, and make few assumptions about the pattern of disturbances (although they implicitly assume an upper limit on their magnitude).

For example the machine you described to etch diffraction gratings (am I remembering this correctly?) must be accurate to within .1 microns, must react in fractions of a second, and probably assumed that earthquakes were rare.

Thus: > >>A control system that can cancel only 10 percent of a disturbance >>isn't much of a control system.

Now consider the squirrel trying to maintain a CEV "I am well fed". This is set by the squirrel's metabolism and diet. On a summer day the squirrel does not need to burn any calories to stay warm (low metabolic load), plus it can find a lot of food easily (good diet). However on a winter day the squirrel must burn a lot of calories to stay warm, and cannot find much food. This is the maximum disturbance case.

If the squirrel cannot make assumptions about the pattern of summer and winter days then it must be able to survive indefinitely in winter. I'm not sure what such a squirrel would look like, but it would be big to keep surface heat losses down. (Maybe a mammoth?)

In reality species (through evolution) have a long history of the climate. Squirrels "know" that they can expect the winter to end (and have a fair estimate of how long winters are). Their bodies and lifestyles are tailored to have enough fat reserves at fall to survive a bad winter, and not much more. If a squirrel stores as much fat as possible (rather than as much as it will likely need) then it won't get around to such crucial things as finding mates to make baby squirrels with.

Many control systems in living organisms make similar assumptions about the pattern of disturbances they will encounter.

I've made a program to illustrate an extremely abstract version of a "loose" control system.

The environment is one dimensional (represented by a number). Time is quantised in ticks. There is a "disturbance" size (default 4), and each tick the environment is changed by this size. (A random number gives a 50% chance of increment, and 50% chance of decrement.)

Three copies of this environment are kept. In one there is no attempt to control it. In the second the number is moved two units towards zero each tick. In the third the environment is incremented by one unit (to introduce a "bias") before being moved towards zero as in the second.

In case that was unclear here is some pseudo code:

```
for (each tick)

    if /* 50% chance of being true */
        change = 4
    else
        change = -4

    free_new = free_old + change
    cont_new = cont_old + change
    bias_new = bias_old + change + 1

    if (cont_new > 0) cont_new = cont_new - 2
    if (cont_new < 0) cont_new = cont_new + 2

    if (bias_new > 0) bias_new = bias_new - 2
    if (bias_new < 0) bias_new = bias_new + 2
```

The tick-by-tick disturbance has a magnitude of 4, however the long term average disturbance is 0 (or +1 in the biased case).

The "gain" of the controlled system is (at most) -0.5, and indeed the error is often non-zero. However it keeps the environment within 8 units of zero 93% of the time.

The biased system has better effective gain if the error is negative, and worse gain if the error is positive. Nonetheless it too keeps the error bounded at small values.

The results of one run (of 10000 ticks) are displayed below. Min and max are the highest and lowest values of the number. Sum is the average of the absolute value of the error, while "err" is the average of the square of the error.

```
%wavers: loops=1 bias=0, incr=4

%wavers: free max= 404 min= -124 sum= 208.36 err=51220.66
%wavers: cont max= 28 min= -30 sum= 3.20 err= 23.29
```

%wavers: bias max= 61 min= -13 sum= 6.67 err= 104.30

The code of the program follows at the end of the post. It generates a graph of the environment in PostScript, as well as the summaries above.

Now for a few comments:


> The importance of the result in the eyes of a biologist doesn't mean
> that there's a control system for the variable. In lower animals, many
> reference levels are inherited and are not adjustable.

I think this point is irrelevant. Your reference level for blood CO2 level is inherited and not adjustable in the same fashion, yet you have a control system for it.

> You'd better draw me a diagram of [the squirrel's] control system.

The concert hall cooler is easier to diagram, and makes the same point. It shows a system that can maintain a CEV (temperature is less than a given value) even though it's instantaneous gain (cooling power) is overwhelmed by the heating power of the audience.

> >We get the same sort of thing in concert halls. The air
> >conditioning cannot handle the heat output of 40,000 sweaty rock
> >fans. [But it] can keep the temperature acceptable even
> >though it can only compensate for a fraction of the momentary
> >disturbance.



Here's a diagram of the system. The top level reference is 35 degrees Celsius (95 F). If the temperature gets higher than that then people start collapsing of heat-stroke, and the hall gets hit with lawsuits.

The climate control ECS uses the input function:

```

delta heat = expected heat from concert * fudge factor -
             cooling capacity of air conditioners
delta temp = delta heat * thermal mass of the concert hall
predicted temp = delta temp + current temp

```

The comparator of the ECS is:

```

if (predicted temp >= max safe temp)
    error = predicted temp - max safe temp

```

```
else
    error = 0
```

And the output function is:

```
if (error not zero)
    turn on air conditioners
```

(Note: this portrays only the cooling side of the climate control system.)

```
> The only control that exists is over a 24-hour period. The average
> temperature is kept within plus or minus x degrees of a reference
> level. That's not very good control.
```

The actual CEV is "temperature is not high enough to be a health hazard", and this is controlled quite well unless the drummer spontaneously combusts.

```
> The only reason this works is that the fans are willing to accept a
> wide range of temperature error, beginning at "too cold" and ending up
> at "too hot".
```

Exactly. Not all control systems are built to microscopic tolerances.

```
> RE: Greg helping Evan to solve a problem, at Evan's request:
>
> >> See how easy it is when nobody is trying to figure out how to
> >>control someone else?
>
> >See why it seems that your definition of control is not:
> > - getting someone else to do what you want
> >but
> > - getting someone to do what they don't want to
>
> It isn't so much getting someone to do what they don't want to as
> getting someone to do what you want, without considering what they
> want. You're controlling in either case, of course.
```

Ack! If you limit "control" in social situations to mean almost instantaneous 100% compliance regardless of the wishes of the controlee then a science of social interactions does not have to consider control any more than it considers unicorns.

If you threaten people with "walk through that door or I'll shoot" you assume that they do not wish to be shot. If you want control regardless you'll have to pick them up and toss them through the door. This passes way outside the realm of psychology, and into the realm of physics.

-----8<-----cut here----->8-----

```
/* a quick-and-dirty program to test a loose control system
*
* This tests three systems that all have the same disturbance.
* They are:
*     free - no control
*     cont - some control
```

```

*           bias - bias in the disturbance
*/

#include     <stdio.h>

#define ABS(x)  (x)>0 ? x : -(x)

#define BINS    512
#define BIN2    (BINS/2)

long   free_bin [BINS];
long   cont_bin [BINS];
long   bias_bin [BINS];

main (argc, argv)

int     argc;
char*   argv[];

{

int     max_loop;           /* time to run           */
int     init_bias;        /* initial value of counters */
int     increment;        /* size of counter increment */
int     d, i, j, k;
int     blip,   last;     /* random booleans       */
int     pos,    neg;
int     up,     down,   high,   low; /* bins to test random numbers */

long   b, max_bin, first_bin;

long   seed;

long   free_tot, free_max, free_min;
long   cont_tot, cont_max, cont_min;
long   bias_tot, bias_max, bias_min;

float  free_err, free_sum;
float  cont_err, cont_sum;
float  bias_err, bias_sum;

printf ("%!PS\n");

if (argc > 1)
    max_loop = atoi (argv [1]);

if (max_loop < 1)
    max_loop = 1;

increment = 4;
if (argc > 2)

```

```

        increment = atoi (argv [2]);

init_bias = 0;
if (argc > 3)
    init_bias = atoi (argv [3]);

printf ("%s: loops=%d bias=%d, incr=%d\n",
        argv [0], max_loop, init_bias, increment);

seed = time (0);
srandom (seed);

printf ("%s: seed=%ld\n", argv [0], seed);

        /* initialize counters and variables */

free_tot = free_max = free_min = init_bias;
free_sum = free_err = 0;
cont_tot = cont_max = cont_min = init_bias;
cont_sum = cont_err = 0;
bias_tot = bias_max = bias_min = init_bias;
bias_sum = bias_err = 0;

pos = neg = up = down = high = low = 0;
last = 0;

for (b=0; b<BINS; b++)
    free_bin [b] = cont_bin [b] = bias_bin [b] = 0;

k = 0;

        /* simulation loop */
for (j=0; j<max_loop; j++) {
    for (i=0; i<10000; i++) {

        blip = random() & 0x100;

        if (blip) {
            d = increment;
            pos++;
            if (last) high++; else up++;
        } else {
            d = -increment;
            neg++;
            if (last) down++; else low++;
        }
        last = blip;

        free_tot += d;
        cont_tot += d;
        bias_tot += d;
    }
}

```

```

        if (cont_tot > 0) cont_tot -= 2;
        if (cont_tot < 0) cont_tot += 2;

        bias_tot++;
        if (bias_tot > 0) bias_tot -= 2;
        if (bias_tot < 0) bias_tot += 2;

        b = free_tot + BIN2; if (b >= 0 && b < BINS) free_bin [b]++;
        b = cont_tot + BIN2; if (b >= 0 && b < BINS) cont_bin [b]++;
        b = bias_tot + BIN2; if (b >= 0 && b < BINS) bias_bin [b]++;

        free_sum += ABS (free_tot);
        cont_sum += ABS (cont_tot);
        bias_sum += ABS (bias_tot);

        free_err += (free_tot * free_tot);
        cont_err += (cont_tot * cont_tot);
        bias_err += (bias_tot * bias_tot);

        if (free_tot > free_max) free_max = free_tot;
        if (free_tot < free_min) free_min = free_tot;
        if (cont_tot > cont_max) cont_max = cont_tot;
        if (cont_tot < cont_min) cont_min = cont_tot;
        if (bias_tot > bias_max) bias_max = bias_tot;
        if (bias_tot < bias_min) bias_min = bias_tot;
/*
        printf ("%s: free max=%5ld min=%5ld sum=%f err=%f \n",
                argv [0], free_max, free_min, free_sum, free_err);
        printf ("%s: cont max=%5ld min=%5ld sum=%f err=%f \n",
                argv [0], cont_max, cont_min, cont_sum, cont_err);
        printf ("%s: bias max=%5ld min=%5ld sum=%f err=%f \n",
                argv [0], bias_max, bias_min, bias_sum, bias_err);
        printf ("%s\n");
*/
    }
    k += i;

    printf ("%s: pos=%ld, neg=%ld, 00=%5d, 01=%5d, 10=%5d, 11=%5d\n",
            argv [0], pos, neg, low, up, down, high);
    printf ("%s: free max=%5ld min=%5ld sum=%8.2f err=%8.2f \n",
            argv [0], free_max, free_min, free_sum/k, free_err/k);
    printf ("%s: cont max=%5ld min=%5ld sum=%8.2f err=%8.2f \n",
            argv [0], cont_max, cont_min, cont_sum/k, cont_err/k);
    printf ("%s: bias max=%5ld min=%5ld sum=%8.2f err=%8.2f \n",
            argv [0], bias_max, bias_min, bias_sum/k, bias_err/k);
    printf ("%s\n");
}

max_bin = 1;
first_bin = 0;
for (b=0; b<BINS; b++) {
    if (free_bin [b] && (!first_bin)) first_bin = b;
    if (free_bin [b] > max_bin) max_bin = free_bin [b];
}

```



```

printf ("\n %% free variables, scale=%ld\n", max_bin);
printf ("[1 1] 0 setdash\n %ld 50 moveto\n", first_bin+50);

for (b=0; b<BINS; b++)
    if (free_bin[b])
        printf ("%ld %ld lineto\n",
            b+50, (long) (600.0 * free_bin [b] / max_bin) + 50);
printf ("stroke\n");

max_bin = 1;
first_bin = 0;
for (b=0; b<BINS; b++) {
    if (cont_bin [b] && (!first_bin)) first_bin = b;
    if (cont_bin [b] > max_bin) max_bin = cont_bin [b];
}

printf ("\n %% cont variables, scale=%ld\n", max_bin);
printf ("[2 1] 0 setdash\n %ld 50 moveto\n", first_bin+50);

for (b=0; b<BINS; b++)
    if (cont_bin[b])
        printf ("%ld %ld lineto\n",
            b+50, (long) (600.0 * cont_bin [b] / max_bin)+50);
printf ("stroke\n");

max_bin = 1;
first_bin = 0;
for (b=0; b<BINS; b++) {
    if (bias_bin [b] && (!first_bin)) first_bin = b;
    if (bias_bin [b] > max_bin) max_bin = bias_bin [b];
}

printf ("\n %% bias variables, scale=%ld\n", max_bin);
printf ("[2 2] 0 setdash\n %ld 50 moveto\n", first_bin+50);

for (b=0; b<BINS; b++)
    if (bias_bin[b])
        printf ("%ld %ld lineto\n",
            b+50, (long) (600.0 * bias_bin [b] / max_bin)+50);
printf ("stroke\n");

printf ("showpage\n");
}
--
De apibus semper dubitandum est - Winni Ille Pu

```

Date: Wed Nov 04, 1992 5:37 pm PST

Subject: Re: description

(penni sibun 921104.1400)

Avery.Andrews (921104.1349)

Rick Marken (921103.1500)

>The important part of the process in the PCT model of language is that
>it is closed loop -- the words that represent the meaning (ws) of what
>is being said are generated by an error signal -- not by the intended
>meaning itself, rml. ...

In a sense, something like this is already built into Salix. The entities and relations in the KB that is to be expressed have a bit which marks `already said', and the generation process is organized so as to go on until everything is so-marked.

this actually isn't quite true. salix (my generation system) doesn't ``say everything.'' when it's told to talk about something, say a house or a family, it's told about a ``completeness criterion''--a specification of what it must make sure to mention. as soon as it's mention the critical things, it can stop. in a house, the completeness criterion is all the rooms (possibly excluding the bathroom; this criterion for house/apartment descriptions is well-documented); in a family, it's something like all the family members still living, or here at thanksgiving. salix starts with a representation of all the objects/properties in a house or family or whatever it's talking about and the relations between them. (actually, salix keeps all its knowledge of everything together.) whenever salix includes an object or a relation in the text it produces, it marks it as already-said. it also can make simple inferences and mark as already-said things that are not explicitly said. for example, if salix says ``penni's father is john,'' both (father-of john penni) and (daughter-of penni john) are marked.

when salix is done w/ a description, it may or may not have marked all the objects as already said. it *never* has marked all the relations it knows about. in principle, it could happen, but it just never does. for instance, in a house, all the rooms bear lots of relations to each other such as proximity (or not), direction (eg, north-of, east-of), and connected via a door (or not). salix's job is to talk about the rooms and how they are related, but it is sufficient to express one or two of those relations. (in fact, in a house description, the direction relations are never directly expressed at all. instead, they are used to dynamically compute whether a room is, eg, ``left of'' or ``right of'' another room. such relations are never marked.)

The problem is that the reference level (the whole KB)

the reference level must be a function of the kb, mediated by the completeness criterion, right?

& the controlled

perceptions (the portion of the KB marked `already said' are not continuous-valued vectors ... -- people just don't know how to make connectionist systems do the sorts of things that speakers do.

exactly. i absolutely agree that my symbolic representations are extremely crude and almost certainly completely wrong. on the other hand, if i were to try to use continuous values, i'd still be trying to figure out how.

[Avery Andrews 921104.1654]

Here's an aspect of the PCT-ification of Salix that I've been overlooking for a while. Salix as it stands simply marks information as said when it thinks it's said it, but what's really going on has to be a bit different. In the first place, for adult conversation,

the things `to-be-said' should really be taken as things `to-be-known-by-the-hearer'; so that their reference `signal' is what the hearer is wanted to know, the perception what the hearer is `perceived' as knowing, where a person's being in the presence of X being said is normally taken as a sufficient basis for concluding that the person knows X (there must be a better way to say this, but it eludes me for the present).

well, remember that there's a little bit of this done by the inference mech., as described above. i agree that this sort of thing needs to be captured somehow. i avoided that sort of thing like the plague when i drew up my thesis proposal, because in ai/comp ling, this sort of thing falls into a black hole called ``user/hearer modeling.' the basic idea of this is that i build a model in my head of what you have in your head, and we pretend infinite regress doesn't happen. anyway, i didn't want to get into that. now i'm not constrained by passing a thesis, i need to look into this in a more intelligent way.

The point of all this: suppose two people, say a mother and father, are recounting some incident (what their kid did) to a third person. They can do this without both of them just reeling off everything they know about the incident, since one person's presenting an aspect of the incident cancels that contributor to the error signal for both of them (I think that what is going on is that one person takes the role of primary narrator, the other filling in with such bits as the first leaves out). Having the loop closed through the environment should make it easier to explain how cooperation works.

right.

With children, I'm not so sure about all this. My impressions are that (a) they just want to say what they want to say (and definitely don't want anyone else to say it) (b) aren't really focussed on conveying new information to a hearer--savoring common knowledge seems to be more like what they're on about.

i do think that some adult communication is like this. as bill suggested and as i hypothesize in my thesis, description (or story telling) is like this, at least in the small. in general, i think lots of language is like this in the small. but interspersed with this, all sorts of interesting other stuff is going on.

cheers. --penni

Date: Wed Nov 04, 1992 6:28 pm PST
Subject: user modelling

[Avery Ahndrews 921105.1313) (penni sibun 921104.1400)

>i avoided that sort of thing like the plague
>when i drew up my thesis proposal, because in ai/comp ling, this sort
>of thing falls into a black hole called ``user/hearer modeling.' the
>basic idea of this is that i build a model in my head of what you have
>in your head, and we pretend infinite regress doesn't happen.

I can't resist posting this little real-life example:

I heard wailings on every side,
and saw no-one who made them;
wherefore, all bewildered, I stopped.

->I believe that he believed that I believed
that all those voices from within the trunks
came from people who were hidden from us.

Therefore the master said: if you break
a little branch from one of those plants
them thoughts that you have will all be cut short.
Dante *_Inferno_* 13.22-30 (Singleton 1970 translation).

The crucial line in Italian is:

cred' io ch' ei credette ch' io credesse
think I that he thought that I thought

clearly recursive just like the English translation.

Avery.Andrews@anu.edu.au

Date: Thu Nov 05, 1992 12:21 pm PST
Subject: dear Rick

from Ed Ford (921105:1030) (From Rick Marken 921104.1400)

>It's over. The 12 year nightmare of greed, division, deceit, and
>selfishness is kaput...

I don't know what election you were watching, my friend, but most of the U.S. Congress is
back, bad check passers and all, the organization that creates and passes the bills and
funds the projects, pork and all. You democrats in control theorists clothing have placed
the problem only on the president. Now it's all yours, completely yours. To quote New
York Sen. Daniel Patrick Moynihan of Bill Clinton's victory. "My God, now it's OUR
deficit."

>ALL people are control systems -- not just rich, protestant,
>heterosexual, white males.

Gee, what do you call Clinton and Gore, Buddhists?

To paid up members of the CSG..the new Closed Loop goes out snail mail tomorrow (Friday,
Nov. 5th), unless there's a glitch at the printers.

Ed Ford

Date: Thu Nov 05, 1992 2:48 pm PST
Subject: RE: dear Rick

From Tom Bourbon (921105 -- 12:31 CST)

Ditto.

Date: Thu Nov 05, 1992 3:26 pm PST
Subject: PCT References, Re: Squirrels, airconditioners

[From Rick Marken (921105.1100)] Jeff Hunter (921104) --

I think it's great that you are trying to model control systems. I would strongly suggest that you try to get hold of Powers' series in Byte (June-Sept, 1979) and my spreadsheet modelling paper; now available on Bill Silvert's listserver but, EVEN BETTER, available in MIND READINGS from CSG Publishing -- just send \$18.00 to G. Williams, 460 Black Lick Rd., Gravel Switch, KY 40328 U.S.A.

In fact, I think this is a good time to remind all of you who REALLY want to learn PCT that there is some VERY IMPORTANT reference material available that should be digested thoroughly before making definitive statements about PCT and how it works. Everyone should, of course, OWN (and re-read periodically) Powers' classic Behavior: The control of Perception. I don't know the ordering info on this but the publisher is Aldine/DeGruyter (which is somewhere in NY). Every serious PCTer (especially would-be modellers) should OWN a copy of every book published by CSG Publishing; there are now four books available from CSG Publishing; Powers' Living Control Systems I and II, Robertson and Powers' Intro to Psychology and Marken's Mind Readings.

Unfortunately, the Powers' Byte articles are not yet in any of these books -- it is important to have copies of these, however.

Now, back to Jeff's post:

> Now consider the squirrel trying to maintain a CEV
>"I am well fed". This is set by the squirrel's metabolism and diet.

A CEV is a variable. It sounds like you are describing the reference for the variable "fed" -- the reference is "well".

> The "gain" of the controlled system is (at most) -0.5, and
>indeed the error is often non-zero. However it keeps the environment
>within 8 units of zero 93% of the time.

You have ignored the dynamics here; this is, indeed, a low gain system but it responds instantaneously and completely to disturbance; it also completes its action before the next disturbance can have any effect. This unusual dynamic (compared to the way variables change in the real world) is what is responsible for the relatively good control exhibited by this low gain system; put such a system in an environment of continuously changing variables (ie. the real world) and add the proper dynamics to the system and you will see that a system with this kind of gain don't control hardly much at all. There are some good tips about computer modelling in the Byte articles.

```
> %wavers: loops=1 bias=0, incr=4
>
> %wavers: free max= 404 min= -124 sum= 208.36 err=51220.66
> %wavers: cont max= 28 min= -30 sum= 3.20 err= 23.29
> %wavers: bias max= 61 min= -13 sum= 6.67 err= 104.30
```

It's excellent to collect data like this. Bravo. But remember, these results apply only for the discreet system (and environment) that is instantiated in you program. Now try a

system that is dynamically realistic -- that STARTS responding to the current effect of the disturbance on each iteration. Dynamics makes a BIG difference that those of us reared on digital (as opposed to analog) computers often tend to forget (like me -- I'm a digital baby, though I do have a dim memory of analog machines).

```
> The climate control ECS uses the input function:
>
> delta heat = expected heat from concert * fudge factor -
>             cooling capacity of air conditioners
> delta temp = delta heat * thermal mass of the concert hall
> predicted temp = delta temp + current temp
>
> The comparator of the ECS is:
>
>         if (predicted temp >= max safe temp)
>             error = predicted temp - max safe temp
>         else
>             error = 0
>
> And the output function is:
>
>         if (error not zero)
>             turn on air conditioners
```

OK, this is a start. Now try to actually run the simulation (you'll need to do quite a bit of work still to make it realistic -- especially putting in those pesky dynamics). But I think that you will get quite an interesting surprise if you set up the model correctly. You will find that your "predicted temp" function doesn't help much (compared to just controlling without prediction) unless you are nice to the control system and make actual variations in temperature (disturbances) a close mimic of the predicted values. But convince yourself; feedforward (which it what it seems you are trying to make a case for -- or predictive control, or whatever you want to call it) doesn't help, unless, of course, the prediction is correct. For some reason, the idea of predictive control is quite seductive; the only way to convince oneself of it's uselessness is through ACCURATE modelling (I use emphasis because with predictive control, it's easy to, well, cheat -- since you can set things up to the benefit of the model).

Anyway, keep up the modelling.

Regards Rick

Date: Thu Nov 05, 1992 3:47 pm PST
Subject: CONNECTIONISM AND hpct

[From Bill Powers (921104.2100)]; Martin Taylor (921104.1340) --

```
>Connectionism depends at root on the distribution of
>representation, not on whether some elements have a continuous
>rather than a discrete range of variation. Your version of HPCT
>isolates the responsibility for the control of particular percepts
>to particular ECSs. We would not do that, but would share the
>responsibility through overlapping quasi-modular groups of ECSs.
>Your version is to connectionist HPCT as classical AI is to
>connectionist classifiers. And it leads to related kinds of problem.
```

The difference you see depends on the dimension you're attending to. For me, the difference between AI and connectionism is that connectionist models let the signals representing variables be the important thing, where in AI all variables had to be converted first into symbols (words, mainly) before they could be operated upon, and then the operations were carried out by rule-driven algorithms for symbol manipulation instead of by computing devices that handle signals directly. To me, that's the difference between analogue and digital computing.

To represent a system with fixed properties, it doesn't matter whether you use a distributed network or an equivalent set of individual functions. BCP, p. 39:

"It is convenient to think of the brain as a collection of localized functions, and of neural signals as occurring in definite pathways linking functions together. The model, however, will not be invalidated if these elements prove some day to be distributed over large volumes of the brain. The organizational properties of this model do not depend on its geometrical properties."

If we want to account for the way these functions come into being, then the network representations will probably be necessary. During maturation, axons grow in ways that depend on what the system is doing, and even after maturation is complete, synapses appear, change their properties, and disappear.

I think that such networks have limited scope; the size of a sensory or motor nucleus. On a larger scale there is clearly an architecture composed of separate modules and separate layers. The types of neurones are different in different modules of the brain; the brain is not just one huge network composed of identical elements operating by identical principles.

The connectionist models I have seen do not impress me as much as they impress their inventors. A great deal of subjective interpretation is involved in saying that a network "classifies" its inputs, or even that it "recognizes" a form. What these networks actually do is a lot simpler than that: given a set of inputs, they will produce certain outputs over a range of the inputs. To label this process "classification" implies first that one already knows what classification is, and second that there is no other kind of perceptual operation of any importance. I reject both implications as unwarranted, the first because there has been no careful investigation of the elements of perception (of the kind I have tried to develop) and the second because it is obvious, at least to me, that a great deal more than classification goes on in perception. I think that connectionists are trying to accomplish in one jump what the real perceptual system does stage by stage. This may be possible, to some extent, but this sort of modeling will necessarily fail to account for perceptions of both lower and higher levels than classifications: perception of motion, for example, or perception of principles.

Best, Bill P.

Date: Thu Nov 05, 1992 4:23 pm PST
Subject: Re: language

(penni sibun 921105.1300)

[From Rick Marken (921104.1400)]

penni sibun (921103.1600) on the relation between her model and pct model

>i think the major difficulty in mapping bet. yr model and mine is the
>``intended meaning'' part.

I agree -- intentions (reference signals) ARE the ONE BIG difference
between PCT and ALL other models of living systems.

nonsense. intentions are one of the many banes of traditional ai.

re: below. i can't make much sense of yr cryptic equations. i don't think you answered
my question. i didn't ask whether you thought there were hidden loops in my model. i
asked how you could make a pct model that accounted for the same phenomena. obviously, i
believe that the phenomena my model captures are important. i'd like to see how you
think pct could account for the same process. could you please try again?

> this is a basic sketch of my model: i cast
>the issue as deciding what to say next.
> another way to characterize this might be 1) where are we in the
>linguistic structure (ws) and 2) where are we in the structure of
>``meaning(s)'' (?) we are talking about.

In order to be able to answer these two questions, doesn't the model
have to be able to perceive the state of ws and "where we are in
the structure of meanings"? That is, 1) and 2) describe perceptual
variables, right?

sure. that's not the issue. i'm not interested in quibbling whether something is a
perception or a perception of a perception.

> ignoring the rest of the
>context for the nonce, what is said next is determined both by 1) and
>by 2). this implies that it is *not* the case that the ``meaning'' is
>selected and then the linguistic structure is found to fit it, and the
>next bit of language is produced.

Your second sentence is not a correct description of the PCT model of
meaning generation as I described it in my diagram; the PCT model
says that a linguistic structure (ws) is generated which moves perceived
meaning into a match with reference (selected or intended) meaning.
You say that 1) and 2) [the questions quoted above] determine what is
said next. To me, it sounds like the perception of 1) where you are in
the linguistic structure (call this p.ws) and 2) where you are in the
structure of meanings (call this p.m) determines what you say next
(call this o). So, your model seems to be $o = f(p.ws, p.m)$ -- definitely
not a control model. But you may have a control model without knowing
it; what you say next may be ws and m rather than o -- so $(p,m) =$
 $f(p.ws, p.m)$ AND $(p.ws, p.m) = g(p,m)$ -- closed loop.

cheers. --penni

Date: Thu Nov 05, 1992 4:30 pm PST
Subject: Loose control systems

[From Bill Powers (921105.1300)] Jeff Hunter (921104)--

>Your contention is that "proper" control systems have loop gains of
>at least -5 (preferably -10 to -100). I believe that this rule-of-
>thumb of yours may be a result of the specs of the systems you
>designed. I also believe that this is too constrained a definition
>to describe the majority of control systems in living organisms.

No, it is a result of the measured parameters of all the control systems we have investigated experimentally. Most such systems have gains well over -30 and some have measured -200. In visual control tasks it is in the thousands -- these control systems seem to use integrating outputs, with an effective steady-state gain too high to measure accurately. Even the E. coli system has a high loop gain, although there is a very large noise component. When a disturbance is applied to an E. coli model, its average value is cancelled perfectly, implying a very high loop gain. Of course it has a low bandwidth; E. coli can't resist disturbances that vary too rapidly.

>A "good" control system keeps its error within a specified distance
>of zero for a specified percentage of the time for a given class of
>disturbances.

You're free to set your own criteria, of course, but in my book this would be a control system that works only intermittently and has no control at all within the "specified distance." In all the examples of ordinary behavior that I know about, control is much better than my relaxed rule of thumb. Driving a car, balancing a checkbook, walking or running, adjusting the contrast on a TV set, baking a cake, typing letters, looking at things, singing, riding a bicycle, drawing a picture, solving an equation, measuring a length, shopping for groceries --- the list is really endless. In all these behaviors, perceptions are controlled so closely that their deviations from the intended states are at the lower limits of detection, despite disturbances of all sorts.

>The control systems you describe generally have a small
>tolerance for error, quick response time, and make few assumptions
>about the pattern of disturbances (although they implicitly assume
>an upper limit on their magnitude).

"Small tolerance for error" is a strange way of putting it. A control system maintains the error as low as it can, given its loop gain. The response time is whatever it is -- some systems are slow and some are fast. Control systems make no assumptions at all -- there is no equipment in them for doing that, and anyway they do not sense "the pattern of disturbances". Their action is based strictly on perceiving the controlled variable itself. They work no differently when the disturbance is caused by invisible events, or when multiple disturbing effects are present, randomly varying, at the same time. If a disturbance that is too large appears, control will simply fail: the output will become as large as it can, and if a larger output is needed, control is lost.

>For example the machine you described to etch diffraction
>gratings (am I remembering this correctly?) must be accurate to
>within .1 microns, must react in fractions of a second, and
>probably assumed that earthquakes were rare.

The sensing of the position of the ruling diamond was accurate to about 0.01 wavelength of laser light, or about 0.006 microns. The machine itself was accurate to about one wavelength, or 0.6 microns. A normal disturbance was corrected in about 1/4 second. If an earthquake had happened, the system would have been able to resist the slow waves but not the fastest ones.

>If the squirrel cannot make assumptions about the pattern
>of summer and winter days then it must be able to survive
>indefinitely in winter. I'm not sure what such a squirrel would
>look like, but it would be big to keep surface heat losses down.
>(Maybe a mammoth?)

Saying that the squirrel "makes assumptions" about such things is an anthropomorphism. If you want to understand how the squirrel works, you have to propose a model of the control systems involved. For example, in warm weather, the squirrel's metabolism is turned down and it needs less food; therefore it gathers less. The reference level for amount of food eaten varies, presumably because it is the output of a more global control system that monitors a number of physiological variables, keeping them near reference levels. When the weather begins getting colder, the squirrel's reference-metabolism is increased, and it gathers more and more food; the control systems actually set the reference level higher than immediately necessary so food accumulates. When the weather is cold enough, the attempt to maintain body weight and other functions at their reference levels begins to fail, and other control systems come into play, preserving the life of the organism by reducing the metabolism to reduce the rate at which reserves are used. This leads to the state of hibernation.

There are many control systems involved in the overall behavior of a squirrel. They are probably all "good" control systems, in my terms. Out of their operation come the superficially-observable actions that we classify as foraging, food-storing, putting on fat, hibernating, and so forth. What we see are only symptoms of the actual control processes and the interactions among control systems.

>In reality species (through evolution) have a long history of the
>climate. Squirrels "know" that they can expect the winter to end
>(and have a fair estimate of how long winters are).

I doubt very much whether any such knowledge resides in a squirrel. Squirrels have evolved to the point where the actions of their control systems are sufficient for most squirrels to survive most winters. There is no need to say that the control systems themselves know anything about winter and summer: a season is much too complex a concept to be perceived by a squirrel. The squirrel can perceive temperature and the state of its own physiology. I think that giving it any more sophisticated abilities is unnecessary and fanciful.

>Many control systems in living organisms make similar
>assumptions about the pattern of disturbances they will encounter.

This is too metaphorical for me. Control systems don't make assumptions about disturbances. They control variables. Where do you put the assumption-maker in a control system diagram?

RE: your program.

>The "gain" of the controlled system is (at most) -0.5, and
>indeed the error is often non-zero. However it keeps the
>environment within 8 units of zero 93% of the time.

That's a meaningless way of evaluating the control system. With a gain of -0.5, the effect of a disturbance on the controlled variable would be reduced to $1/(1 + 0.5)$ or $2/3$ of the value it would have had with no control at all. So $2/3$ of the effect of the

disturbance goes uncorrected. If that's good enough for some purpose, then there's not much point in using any control system at all. You'd only gain 33% over a system that just took the effects and suffered them. I know you will respond, "But maybe that 33% makes all the difference!" If so, OK. But you still don't have a system with much ability to control what is happening to it. It's just lucky that control this poor suffices. That is not true for most of the variables that organisms have to control.

I haven't run your program yet, but have no reason to doubt your description of its behavior.

I think you badly underestimate the quality and number of the control systems that are actually involved in real behavior. Most variables that are of importance to organisms are maintained very accurately at their reference levels under all normal conditions. For any organism there are conditions that strain the control processes and even overwhelm them; when that is the case the organism just suffers the consequences. It can do no better than its best. When things get bad enough it dies.

You are imagining an environment in which there are only slight fluctuations in conditions, so that only a slight capacity for control is enough to assure survival, or where the system is pushed to its limits, and a small increase in control capacity will have critical effects. Such situations might exist. But by focusing on them, you ignore the vast array of circumstances in which control must be and is precise and quick. You ignore, in fact, all the major phenomena of behavior, which your model would be totally incapable of reproducing. Phenomena such as the way a squirrel climbs a tree and jumps from limb to limb; the way it spots food and finds its way to it and picks it up and carries it back to its nest; the way the squirrel evades a chasing cat and scolds a pestering bird; the way the squirrel simply stands up and walks. By looking at fringe phenomena and tiny effects, you are missing all the big obvious foreground effects that need explaining and that only control theory can explain.

A loose control system controls loosely. Of course. But organisms don't.

Best, Bill P.

Date: Thu Nov 05, 1992 4:49 pm PST

Subject: Re: dear Rick

>>It's over. The 12 year nightmare of greed, division, deceit, and
>>selfishness is kaput...

>

>I don't know what election you were watching, my friend, but most of
>the U.S. Congress is back, bad check passers and all, the organization
>that creates and passes the bills and funds the projects, pork and all.
>You democrats in control theorists clothing have placed the problem
>only on the president. Now it's all yours, completely yours. To quote
>New York Sen. Daniel Patrick Moynihan of Bill Clinton's victory. "My
>God, now it's OUR deficit."

>

>>ALL people are control systems -- not just rich, protestant,
>>heterosexual, white males.

>

>Gee, what do you call Clinton and Gore, Buddhists?

Tom Bourbon added:

>Ditto

Ouch. Et Tu, Tom?

Well, let me just try to correct this little disturbance (for the sake of my own perceptions, of course).

First, a quick note on "bad check passers"; it might sound great if you already dislike congress people but these people were hardly abusing the system; I don't even fault the republicans who did it. These folks were led to understand that they could write checks against their account even though their paycheck was not yet deposited. They were taking a perk that, perhaps, was not explicitly provided, but this was certainly not craven abuse of power; just a misunderstanding (except in a very small number of cases). It was a non-issue and the election results suggest that it was not much of a factor.

As for the congress being the source of all our ills -- well, I agree that there is a great deal of "pork barrel" stuff going on, but it is rather bipartisan. And the president is supposed to set the agenda, right? And there have been presidents who have been able to successfully deal with congresses of the other party. It ain't easy to deal with the special interests in congress -- but that's part of the job of the president, I'm afraid. Maybe he should have tried a little cooperation rather than confrontation.

But the main reason for my remarks was this:

My image of a president is of someone who should lead by articulating as best as possible a consensual system concept which, if adopted by most people in the community, could lead to cooperative efforts that improve individual control. Reagan and Bush did not articulate (to me) such a system concept (well, Reagan kind of did). My own perception of the system concept communicated during this era was one that encouraged interpersonal conflict and greed as an economic policy, that was actively hostile to cultural diversity, and that encouraged only one approach to conflict resolution -- force. All this is just perception but I felt like a definite outsider during this whole period (Bush even said that I should be disenfranchised -- he said that athiests are not really american citizens; silly and harmless but it contributes to a perception). So there is a certain zeitgeist that is communicated by leadership; and it is that zeitgeist that I found very uncomfortable (to me -- that's just a reflection of my reference for a system concept; apparently, 37% of the people were at least OK with the Reagan/Bush zeitgeist).

I think the loop gain between actions of president/congress and individuals in society is pretty loose. But there are some influences; not always intentional. Something did happen during Reagan/ Bush that is upsetting to me: it came to pass in that time that the wealth controlled by the top 1% went WAY UP and that controlled by the lowest 50% went WAY DOWN. Now I don't think Reagan/Bush tried to make this happen ON PURPOSE and I don't think they influenced this on their own (congress is part of the problem) but it happened; and it is a result that is consistent with the system concept I was picking up: "conflict is good and the strong survive and get the goodies. The poor, the weak, etc, ie the people who didn't vote, are just left behind -- the fact that they were out on the streets is just evidence of their incompetence". Well, that's their system concept. I'm happy because I see a new president who is already starting to articulate a cooperative system concept, a non- punitive system concept. It's not perfect (from my point of view) but its a hell of a lot closer to the kind of system concept I like than was the one articulated by Reagan/Bush.

Of course, I also like Slick because 1) that was my nickname too, when I was a kid and 2) he is exactly the same age as I am. We went through the same stuff; we experienced the Vietnam era from the same perspective, we're both very smart and cute, we both have grey hair, we both have very smart/very cute wives; we both value knowledge over belief; we both care about people; we are now in charge.

I now return you to your regularly scheduled control theory discussion.

Regards Rick

Date: Thu Nov 05, 1992 5:09 pm PST
Subject: Re: language

[From Rick Marken (921105.1400)]

In article <92Nov5.124237pst.29192@hmmm.parc.xerox.com> Penni Sibun
<sibun@PARC.XEROX.COM> writes:

>(penni sibun 921105.1300)

>

> [From Rick Marken (921104.1400)]

>

> penni sibun (921103.1600) on the relation between her model and pct model

>

> >i think the major difficulty in mapping bet. yr model and mine is the

> >`intended meaning' part.

>

> I agree -- intentions (reference signals) ARE the ONE BIG difference

> between PCT and ALL other models of living systems.

>

>nonsense. intentions are one of the many banes of traditional ai.

Well, maybe you are right. They do talk about intentions a lot. My impression is that in most cases the intention is modelled as a command to produce action rather than as a specification for perception. But I'm sure that my impression is wrong (from your point of view) -- perhaps you could describe such an ai model.

>re: below. i can't make much sense of yr cryptic equations. i don't
>think you answered my question. i didn't ask whether you thought
>there were hidden loops in my model. i asked how you could make a pct
>model that accounted for the same phenomena. obviously, i believe
>that the phenomena my model captures are important. i'd like to see
>how you think pct could account for the same process. could you
>please try again?

Ok. This is a very fair question. Perhaps you could describe the phenomenon that your model "captures". Maybe give some quantitative data on the phenomenon. From your discussion with Avery it sounds like your model (Salex, right?) produces sentences that describe scenes. Some of the descriptions are based on inferences derived from the scene but not explicitly "stated" -- these are verbally described scenes, right. I think you already have a discreet PCT model -- which is fine. The references are each of the "descriptions" that should be produced. The perceptions are logical variables -- either they are true or false. The references for each variable is set at true. Then you go through your inference engine and parser and produce states of those variables. A higher

order system is looking for all the reference perceptions to be set to true -- then you are done. So you probably already have a control model -- controlling logical variables.

But I think it would be good to see the data on the phenomenon before judging the value of the model -- control model or not.

Best Rick

Date: Thu Nov 05, 1992 5:16 pm PST
Subject: Jeff's Loose control system

[From Bill Powers (921105.1430)] Jeff Hunter (921104) --

I've got your program to run, and think I understand it now. By the way, Turbo C doesn't like long integers as indexes to arrays -- I fixed that, plus the difference in random() functions.

Your control systems are contained entirely in the lines

```
if (cont_tot > 0) cont_tot -= 2;
if (cont_tot < 0) cont_tot += 2;

bias_tot++;
if (bias_tot > 0) bias_tot -= 2;
if (bias_tot < 0) bias_tot += 2;
```

The perceptual function and comparator are in "if (cont_tot > 0)", for that implies perception of the value of cont_tot and comparison with the reference level of 0. The error signal is either 1 or 0, so this is a binary comparator. The output function either adds 2 to the environmental variable cont_tot (or bias_tot) or subtracts 2. The output function, too, is binary, with a gain of 4 and an offset of 2.

The environmental part of the loop is a simple integrator. If you add 2 on each iteration the environmental variable *_tot increases by 2 and holds its value until the next iteration. The environmental variable and the perceptual variable are the same, here. The output quantity has the values 2 or -2, so the system as a whole can never come to a steady state (with all variables unchanging).

The steady-state loop gain of these control systems is nominally infinite. The reason is the integration in the environmental part of the loop. Loop gain, as you know, is found by multiplying together all the steady-state multiplying factors in the functions encountered in one trip all the way around the loop. The input function gain is 1, because the perceptual signal is the same as the value of the environmental variable. The gain of the comparator is -1, and of the output function is 4. The gain of the environmental feedback link is, in the steady state, infinite, because of the integration. The loop gain is therefore $1 * -1 * 4 * \text{infinity}$, or infinity. So this system would bring its controlled variable exactly to the reference value of 0, if the output were continuous instead of bang-bang. The size of the output quantity determines how fast the error is corrected. In your design, it also determines how much oscillation there is at asymptote.

The control systems reduce the (squared) error that would be seen without control (free_cont) from 51220 to either 23 (unbiased) or 104 (biased), by your numbers. These are sums of squared errors, so the RMS error over 10,000 iterations is 5, 0.0023, and

0.0104. (That value of 5 seems like too great an RMS error for equal distribution between positive and negative disturbances. Better check that). It is impossible to estimate the effective loop gain except for the biased case, where a constant of 1 was added on each iteration. Over 10,000 iterations, this should have added 10,000 to the value of bias_tot (and a lot more to the squared error); instead, bias_tot changed by only 104, or about 1% of the expected value. This yields an effective loop gain of about 100. The effective loop gain is not infinite because of the binary comparison and output functions.

This is therefore not a loose control system but one tight enough to exceed by a factor of 10 my definition of a "good" control system (whether biased or not), with the defect that it must continually oscillate around its final average state. It is the oscillation that prevents the effective loop gain from being infinite and the final error from being zero.

If you intended to portray a loose (low-loop-gain) control system, you didn't succeed. By choosing an environmental variable that would accumulate the effects of outputs and hold them indefinitely, you made the loop gain infinite. I define loose control as the result of a negative loop gain considerably smaller than 10. Perhaps you have been confusing the output sensitivity with the loop gain. In many real control systems, it is true that most of the loop gain is determined by the output function. But in this case, the environment contributes most of it.

Best, Bill P.

Date: Thu Nov 05, 1992 5:55 pm PST
Subject: Re: Dropping Babies

[From Dick Robertson 921105]

Gary, i got this example from a real event. A long time ago someone was telling
What do the rest of you think in answer to Gary's question?

On another note- I still aim to reply to Rick Marken from a week or two back,

Cheers to all, Dick

Date: Thu Nov 05, 1992 6:10 pm PST
Subject: Re: dear Rick

How much does it cost to be a paid-up member of the CSG?

Eileen

Date: Thu Nov 05, 1992 6:53 pm PST
Subject: language, error signals, equilibrium

[Avery Andrews 921106.1348]

Aside from the one Martin Taylor pointed, out, another difference between language production and ordinary control situations stems from the discreteness inherent in language. One can turn, say, Salix, into an environmentally closed loop system by marking bits of the knowledge base as said when they are heard and understood as said by

the parsing system (not part of Salix at this point, but anybody can write a parser these days).

But then the discrete nature of the KB has an interesting consequence. The error signal that drives production is the difference between what is to be said and what has been said. But (a) once something has been said it stays said (b) you can, in many cases, say it all. So rather than having the typical case be an equilibrium state where the error signal is just enough to maintain the controlled quantity next to the reference level, which is thought of as constant, the typical case is that the `reference level' changes (e.g. some stuff becomes wanted to be said), then out comes behavior to get it said, and then activity stops (conversations really do just grind to a halt sometimes). E.g. the whole thing is really much more like an S-R situation, in spite of not really being one.

Avery.Andrews@anu.edu.au

Date: Thu Nov 05, 1992 7:06 pm PST
Subject: CSG Membership

from Ed Ford (921105:2005)

To Eileen Prince and other non-CSG members.....

Membership in the Control Systems Group is \$45. This supports the Closed Loop, our quarterly edition of a condensed version of the CSGnet. It also allows us to give scholarship money to students who would otherwise not be able to attend our annual conference in Durango, scheduled in 1993 from Wednesday evening, July 28th through Sunday morning, August 1st. Those signing up midyear get the back issues of Closed Loop for that year. Part of the conference fee is your dues, so that's when our new year begins. If you are interested in joining, send \$45 payable to Control Systems Group and mail your check to Mary Powers, 73 Ridge Rd., CR 510, Durango, CO 81301.

Best, Ed...

Date: Fri Nov 06, 1992 6:49 am PST
Subject: Re: dear Rick

I am just a lowly grad student who has just signed-on to this list primarily as an observer, but in reply to Rick Marken's last post re: presidential control systems:

DITTO

Regards, M. Lee

Date: Fri Nov 06, 1992 8:35 am PST
Subject: Re: dear Rick

[From Rick Marken (921106.0800)]

> DITTO
>M. Lee

Thanks. You made my morning.

You are most wise and most welcome to the list.

How about a quick autobio and/or a comment about what got you onto csg-1?

Best Rick

Date: Fri Nov 06, 1992 8:37 am PST
Subject: Re: CONNECTIONISM AND hpct

[Martin Taylor 921106 10:45] (Bill Powers 921104.2100)

> The difference you see depends on the dimension you're attending to.
> For me, the difference between AI and connectionism is that
> connectionist models let the signals representing variables be the
> important thing, where in AI all variables had to be converted first
> into symbols (words, mainly) before they could be operated upon, and
> then the operations were carried out by rule-driven algorithms for
> symbol manipulation instead of by computing devices that handle
> signals directly. To me, that's the difference between analogue and
> digital computing.

Fair enough. There are certainly two dimensions (at least) in the comparison. But it appears to me that HPCT allows algorithmic manipulations at the program level and higher (possibly as low as the category level), so the analogue-discrete distinction doesn't fully distinguish PCT from AI.

> "It is convenient to think of the brain as a collection of localized
> functions, and of neural signals as occurring in definite pathways
> linking functions together. The model, however, will not be
> invalidated if these elements prove some day to be distributed over
> large volumes of the brain. The organizational properties of this
> model do not depend on its geometrical properties." BCP, p. 39

I have a new copy of BCP in its plastic wrapping on my desk, but I haven't read it for over a year now, so I didn't remember this paragraph. However, I would argue that one could go further: not only is the model not invalidated, it is enhanced by distribution. Indeed, I think distribution is the only way that a full-blooded PCT system could work in a real, as opposed to a modelled world. A non-distributed (localized, near-orthogonal ECS) system will be as brittle as a rule-based AI system (perhaps I exaggerate) as compared to a neural network performing a similar task.

> The connectionist models I have seen do not impress me as much as
> they impress their inventors. A great deal of subjective
> interpretation is involved in saying that a network "classifies" its
> inputs, or even that it "recognizes" a form. What these networks
> actually do is a lot simpler than that: given a set of inputs, they
> will produce certain outputs over a range of the inputs. To label
> this process "classification" implies first that one already knows what
> classification is, and second that there is no other kind of
> perceptual operation of any importance. I reject both implications as
> unwarranted, the first because there has been no careful
> investigation of the elements of perception (of the kind I have tried
> to develop) and the second because it is obvious, at least to me,
> that a great deal more than classification goes on in perception. I
> think that connectionists are trying to accomplish in one jump what

> the real perceptual system does stage by stage. This may be possible,
> to some extent, but this sort of modeling will necessarily fail to
> account for perceptions of both lower and higher levels than
> classifications: perception of motion, for example, or perception of
> principles.

Here, I fully agree. But the performance of the networks on the tasks set them is impressive to more than their inventors. It impresses me. What I think (and thought for many years before I heard of PCT--I did my first try at neural nets in the early 60s) is that neural networks are totally misused when they are limited to classification. They are good at that (by which I mean the discrimination among predetermined complex patterns in space and time) but they are better as transformers and coordinators. I have for a long time thought that the appropriate use of NN technology was to provide ways of linking stimuli to behaviour, and since learning of PCT it seems to me clear how this can be done effectively.

I think I'll leave to another posting with a different focus the importance of the "squashing function" in the perceptual input path. It came up last year, and was apparently misunderstood. It relates directly to the question of distribution, but needs some technical preparation.

Martin

Date: Fri Nov 06, 1992 9:38 am PST
Subject: Re: dear Rick

From Tom Bourbon (921106 11:01 CST)-----

(Rick Marken 921105 15:47) "Ouch. Et Tu, Tom?"
Re: my "ditto" to Ed Ford's remarks, I believe Ed meant that he would like to call down a pox on both parties and on both the congress and past presidents. (Is that it, Ed?) I agree with him. I could say more, but this is not the place. After hours at the next meeting of CSG?

Hasta luego, Tom Bourbon

Date: Fri Nov 06, 1992 12:33 pm PST
Subject: language issues

[From: Bruce Nevin (Mon 92112 07:38:16 through Fri 92116 14:05:06)]

Awfully hard getting any consecutive time to respond to all the interesting stuff going by.

(Avery Andrews 921026.0928) --

>I'd agree that this little story of mine would stand or fall on the
>basis of careful examination of this kind of case. But if this story
>falls, we're probably left with something like GB as the best bet for
>a theory of grammar!!.

Are the choices really so few and so constrained? This seems based upon an argument of the following form:

Either A or B.
If A, then something like GB
If not (this little story of mine) then not B

Here, A is something like the claim that "learners don't have access to negative evidence," i.e. they are not provided with starred examples, and B is the claim that they are.

For non-linguists: in Generativist writings one puts an asterisk before examples that "you can't say" (at least according to the author). Hence, these are "starred" examples. (I used to call them "asterisky.") For example, the linguist might consider sets of roughly equivalent sentences like this:

- 1a John gave the Ramadan bake sale money to the fund.
- b John gave to the fund the Ramadan bake sale money.
- c John gave the fund the Ramadan bake sale money.

On the basis of this and other sets of sentences, the linguist might postulate a rule that says you can move a "dative" phrase--"to x" after a verb with some kind of "give" meaning--to a position before the verb, and delete the preposition "to".

This rule predicts that the following sentences are also equivalent to each other:

- 2a John donated the Ramadan bake sale money to the fund.
- b John donated to the fund the Ramadan bake sale money.
- c *John donated the fund the Ramadan bake sale money.

However, for many English speakers (including Avery), (2c) is not acceptable, hence the asterisk or star. Others, focussing on the meaning of (2c), may overlook the syntactic or stylistic anomaly; Bill is evidently among these, but it is possible that he represents a third set of others for whom (2c) is not anomalous at all.

Returning to the issue at hand: the claim that language learners do not have access to negative evidence seems to amount to a claim that they do not experience error when they produce syntactically anomalous utterances. The claim is usually that overt correction ("don't say X, say Y") is such a miniscule part of the language learner's experience, and covers so few of the possible errors, that it alone cannot account for the incredible amount of learning that takes place in such a short time (hence much of the structure of language cannot be learned at all and instead must be hard wired in the genome, etc, etc).

Several assumptions underly this claim.

Doesn't this assume that overt correction by other speakers is the only source of negative evidence? It should be obvious to students of PCT that there are many occasions for experiencing error aside from attempts of others at social control. For starters, consider the experience of not being understood and having to try again, or the experience of having the other paraphrase what one has said, either in confirmation or in replaying it to someone who did not understand clearly, both of which are *very* common experiences for children learning a language.

Generativists assume that the learner is faced with a range of alternative grammars, and that learning is a process of eliminating those that don't work. On this view of how learning works, counterexamples have an enormously important and pervasive role. GB

(Government-[and-]Binding Theory, now being supplanted by Chomsky's latest "minimalist" proposals) carries this to an extreme. UG (Universal Grammar), hard wired into the genome, provides for all possible languages. Learning one particular language involves setting parameters in UG so that one possibility is admitted for that parameter and the others are precluded. An example might be the choice of SOV (subject object verb) word order, as opposed to the alternatives. Having made that choice, certain other choices ripple out as entailed consequences (order of modifiers relative to modified words, etc.). Much ink has been spilled in a search for the minimal parameters whose settings cannot be predicted from the settings of other parameters, or which once set can be the basis for predicting a larger number of others, etc. (In this research climate, language differences and variation are not always so closely investigated as they might be, but that can be only a secondary side comment here.)

Generativists also commonly perceive that language structure is so very complex and arcane as to defy a child learning it at all absent biological inheritance of much of the complexity in UG. It is true that Generativist descriptions of language are very complex and arcane. It is also true that not all ways of describing language are so.

None of these assumptions is very convincing to me, but they do keep a number of people employed as linguists (and a number of others not employed as linguists). It also puts a great many properties of language in genetically determined UG and therefore outside the realm of perceptual control.

Returning to the examples, sentence (2c) is somewhat peculiar for me, but not completely unacceptable, so I mark it with a question mark rather than an asterisk:

2c ?John donated the fund the Ramadan bake sale money.

Contrast (2c) with an example that unequivocally merits an asterisk:

3. *John fund the donated.

In operator grammar, the "to" may be zeroed with a verb after which it is so strongly expected as to be redundant. (In some UK dialects, it need not be immediately after the verb, e.g. "John gave it me." Similarly in Danish.) Thus, the reason given in Operator Grammar that (2c) is peculiar is that speakers of English do not have a sufficiently strong expectation of the preposition "to" occurring after "donate."

One might suppose this is because "donate to" simply is not used as frequently as "give to." It is true that the dative reduction is more or less peculiar for many other verbs in the "give" set. Call them the "donate" subset, including:

- ? John administered [to] the fund the money
- ? John allocated [to] the fund the money
- ? John apportioned [to] the fund the money
- ? John communicated [to] the fund the money
- ? John consigned [to] the fund the money
- ? John contributed [to] the fund the money
- ? John conveyed [to] the fund the money
- ? John dealt [to] the fund the money
- ? John delivered [to] the fund the money
- ? John dispensed [to] the fund the money
- ? John distributed [to] the fund the money

? John entrusted [to] the fund the money
? John purveyed [to] the fund the money
? John relinquished [to] the fund the money
? John rendered [to] the fund the money
? John surrendered [to] the fund the money
? John transferred [to] the fund the money
? John transmitted [to] the fund the money
? John vouchsafed [to] the fund the money

But a considerable number of other verbs in the "give" set are acceptable with the transposed argument order and zeroed "to". Call them the "give" subset:

John accorded [to] the fund the money
John allotted [to] the fund the money
John assigned [to] the fund the money
John awarded [to] the fund the money
John conceded [to] the fund the money
John furnished [to] the fund the money
John granted [to] the fund the money
John handed [to] the fund the money
John left [to] the fund the money
John lent [to] the fund the money
John offered [to] the fund the money
John paid [to] the fund the money
John presented [to] the fund the money
John proffered [to] the fund the money
John provided [to] the fund the money
John sent [to] the fund the money
John supplied [to] the fund the money
John yielded [to] the fund the money

I don't think a "frequency of occurrence" argument holds up.

Now, according to Bill's intuitions about how language works, we use words because they are associated with perceptions that we are trying to get our audience to attend to (in the environment or in imagination). Examining these two lists, I don't see any obvious differences in some sort of "to" perception associated with one more strongly than the other.

Instead, the difference appears to be a pretty arbitrary, learned, socially shared perception that "to" is redundant after verbs in the "give" list, but not after those of the "donate" list. Maybe someone else can pull out some generalization that I'm missing.

There is another factor, one that motivated my putting the modifier "Ramadan bake sale" in the examples of (1) and (2), and that is that by minimizing the distance between an operator and its arguments (distance in terms of intervening words that could also be candidates as arguments), we reduce the load on short-term memory for perception of operator-argument dependencies. Thus, when we have alternative orderings for the arguments of an operator we prefer to put the shortest one first (counting all its modifiers):

4a John handed the money over to the fund.
b ?John handed over to the fund the money.

Sentence (4a) is definitely better than the transposed version (4b). However, watch what happens when you lengthen "the money" with some modifiers:

- 4a John handed the Ramadan bake sale money that Kathy had collected from the first half of her list of groups over to the fund.
- b John handed over to the fund the Ramadan bake sale money that Kathy had collected from the first half of her list of groups

Now consider the original two sets of sentences without the modifier:

- 1a' John gave the money to the fund.
 - b' ?John gave to the fund the money.
 - c' John gave the fund the money.
-
- 2a' John donated the money to the fund.
 - b' ?John donated to the fund the money.
 - c' ??John donated the fund the money.

Notice that (1b') is a bit awkward but (1b) is more acceptable:

- 1b' ?John gave to the fund the money.
- 1b John gave to the fund the Ramadan bake sale money.

This is seen in the "donate" sentences as well:

- 2b' ?John donated to the fund the money.
- 2b John donated to the fund the Ramadan bake sale money.

The fronting of the second argument of "give" (with its argument- indicator "to") is motivated by its length with its modifiers in (2b), but there is no such motivation in (2b').

Control for minimizing the distance between operators and their arguments is presumably universal, but to say it was a property of Universal Grammar would be fatuous. It is surely a universal property of control systems that are controlling for sequences that may be interrupted by other controlled sequences. So likewise, I believe, many other claims for UG are really claims for the universality of control.

(Bill Powers (921024.0830)) --

> Avery Andrews (921024.1519) --

> >So the non-occurrence of the second (a piece of `negative evidence') becomes accessible to the learning system.

> Control theory handles the non-occurrence in terms of a reference signal that demands the occurrence. When there is a reference signal specifying that some perception occur, then as soon as the reference signal is set there is an error, which persists until the perception occurs. In this case it is not occurrence of the sentence that matters, but of the meaning.

Maybe you've sorted it out by now. You and Avery were operating at right angles to each other throughout much of this exchange. In this instance, "non-occurrence" meant for Avery a starred example--something marked as not acceptable or not occurring in

acceptable usage, or perhaps (by the reification that is normal in Generativist literature) not occurring in the grammar. At issue above is the question who if anyone other than the linguist has a reference signal for such a thing occurring; and it is not the meaning of the sentence that matters, but whether or not you can express the given meaning in that particular way.

(Avery Andrews (921024.1519)) --

>A further prediction is that the kind of optionality above will be
>an inherently unstable feature of languages: if two such forms are
>used with no discernable difference in meaning, the language-
>acquisition systems of the speakers will be constantly reorganizing
>without being able to find an error-free configuration, &
>presumably at some point one of the forms will other drop out, or
>they will acquire subtly different meanings

Something like this was advanced by Kurylowicz, re the classic observation that "doublets" (specifically an inherited but "irregular" form alongside the form used by an innovating group, not necessarily a younger generation, that used analogy to create a more "regular" form) tend to become differentiated semantically. I think this is part of a more general process whereby distinctions tend to be exploited as differences that make a difference, or else are no longer maintained. But this is a *social* process (i.e. ongoing negotiation of agreements about what a given distinction either constitutes or means). Your formulation here does not differentiate between this social process and the processes of individual language learners, hence Bill's inference:

>The implication is that the mature speaker will come to express the
>same meanings in the same words all of the time, never paraphrasing or
>varying the wording. I'm not sure you would want to maintain that.

What prevents rapid closure, smoothing out all the bumps and ripples in the process of learning a language, is precisely the fact that it is not a process that any one person can carry out in isolation, but rather a social negotiation of agreements, and the prior commitments and motivations of the participants are not necessarily commensurate with each other. Pronouncing "bird" as "boyd" constitutes a token of membership in a community that talks that way, and presenting oneself to others by using an ensemble of such tokens means different things to different people, and on different occasions. A person who can and does pronounce "bird" both ways may in fact be motivated NOT to smooth out the difference. (Use one in Brooklyn, use the other in my job as a newscaster.)

As you say,

>if two such forms are
>used with no discernable difference in meaning

But this is to say that they are used without ever experiencing error, no matter which is used. You suggest that the availability of alternative ways A and B of expressing meaning M, where A and B are not distinguished from one another (no A-circumstance where it is an error to use B, and vice versa), is itself an occasion for error. But this is a sort of meta-error, of a different logical type from the the first that you proposed, as I understand it:

1. A person hears utterance x
2. The person's "comprehension system" comes up with meaning m.
3. The person's "production system" generates in imagination a set of

utterances $U = \{w, y, z\}$, where each member of U expresses meaning m .
4. The absence of utterance x from U is occasion of an error signal.

Apparently, the person guesses at the meaning m . There must have been some basis for this guess--word choice, context, part of x is like part of w , part is like y , and so on. It seems to me that what one typically does with this sort of error is to perceive the other person's utterance x as an error of production or reception or both, perhaps as a repair or a change of syntactic horses in mid-sentence, etc. One's conversational response (or one's strategy as a listener or reader) includes heightened attention to those elements (words, word dependencies) about which one is less certain. One attends also to some alternative meaning m' , and its alternative utterance set $\{p, q, r\}$ also generated in imagination in parallel. One of these may suddenly pop up to most favored status. Socalled "garden path sentences" provide simple examples of this process: "The horse raced past the barn fell down".

In the meta-error case, the person wishes to express meaning m and finds that there are alternative ways $\{w, y, z\}$ for expressing it. If there is some motivation for choosing one rather than another, some difference that makes a difference to the person, then the choice is obvious and there is no error. Commonly, it seems to me, the unthinking preference is a repetition of words or of a syntactic construction recently used in the same discourse (conversation, speech, text, etc.). A more polished speaker may make a different choice with deliberate artifice, perhaps for emphasis or simply for variety. Again, so long as there is a basis for choice, I don't see any occasion for error.

Error may occur for example when choice w vs. choice y constitutes a token of social membership (like the different pronunciations of "tomato," "neither," and "economics"). This is not a difference of meaning in the sense that Bill intends, I think. (That's why I say it "constitutes" instead of "means" social membership. It's performative, in the same sense that "I hereby declare this a disaster area", spoken by a duly empowered official, constitutes a change to the status that the phrase "disaster area" by itself means.)

You (Avery) had specifically in mind examples like (2a') and (2c'):

- 2a' John donated the money to the fund.
- 2c' ?John donated the fund the money.

You propose that just because the learner's "production system" at time t produces both of these indistinguishably to express a given meaning, then every time the learner hears (2a') an error signal results. The language learner's response, you proposed, is to eliminate that error by reorganizing, specifically, by altering the production system so that it no longer produces (2c').

I think that a one-step elimination, analogous to marking (2c') with an asterisk (or question mark), will not do. Such a mechanism must surely function to eliminate production of (1a') every time the learner hears (1c') and to eliminate production of (1c') every time the learner hears (1a'):

- 1a' John gave the money to the fund.
- 1c' John gave the fund the money.

Instead, perhaps there is an incremental change in the expectability of one form or another. Since there may be three or four or more alternative forms in some cases, the best move it would seem to me is in the opposite direction: to *raise* the expectability

of the form that actually occurs, rather than lowering that of the ones that do not occur. But no error signal is needed for this, and no reorganization.

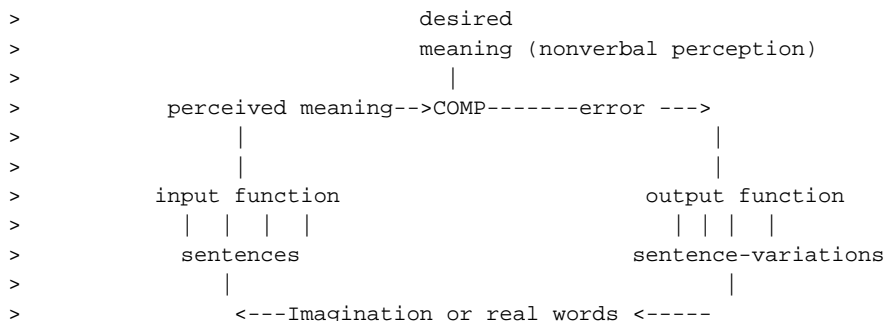
In fact, I see no obvious reason to suppose that a construction such as (2c') could not hang around for the life of the language learner--one of many constructions predicted on analogy to forms that the learner has heard ("predicted by the rules of the grammar" in the customary hypostasis), but not itself actually encountered. I think that it would be untenable to claim otherwise in the face of language novelty, language change, and our capacity for understanding the meaning of utterances across dialect variation, non-native usage, speaker error, difficulties of hearing, and so on.

This interpretation accords with the fact that the acceptability of sentences is not a binary, yes/no property, but rather a graded property, and probably complexly so (that is, on more than one graded parameter in parallel, e.g. differentiated by subject-matter domain for starters).

So I think there are other ways of showing that language acquisition, like all control processes, is guided by perceptual error ("negative evidence"), and that the alternatives are not limited to your proposal and GB.

==+==

(Bill Powers (921025.1900)) --



> Going up the left side, you perceive sentences, which then give rise
> through some sort of input function to perceived meanings. The
> perceived meanings (which are not in words, but in terms of
> perceptions of the things that the words mean) are compared directly
> with the desired meanings. The error acts through an output function
> to vary the construction of sentences. These are the sentences that
> are perceived, closing the loop.

This diagram makes the point about arriving at outputs by continuous control of perception rather than by top-down realization of some sort of archetype. The actual flow of control is probably quite different.

Penni asks (penni sibun 921026):

```

>where are the symbol structures coming from? why are any necessary?
>what does meaning have to do with anything? why can't speech actions
>be just like other actions, which presumably aren't mediated by
>`meaning'?

```

I would suggest that all perceptions "have" meaning, that is, any given perception is linked in associative memory to other perceptions. I'm not sure precisely what you (penni) mean by the phrase "mediated by meaning". Let me lay out a bit of what I think is going on.

Even before we consider meanings, language perceptions involve most of the proposed levels of the perceptual hierarchy, if not all. Phonology alone runs from intensity level up to at least the event level, possibly the program and even principle levels if some ways of doing phonology survive the eventual PCT shakeout. Morphemes and words are event perceptions in Bill's usual estimate of things. (Bill: Is this problematic if syllables or semisyllables turn out to be events? If so, word events would be constituted of syllable events.)

Nonverbal perceptions obviously involve all the proposed levels of the hierarchy, independently and in parallel to these language perceptions.

Event-level word perceptions are linked with nonverbal perceptions. Bill has suggested that the input of a category perception ECS can be satisfied by various nonverbal perceptions (glimpse wagging tail, barking sound ==> dog) or by a word ("dog") or both. I have suggested that the link is in associative memory, and that it is not a simple 1-1 correlation of words and categories.

Nonverbal Perceptions		Word Events
P1 P1' P1"	<---->	W1
P1' P16 P17	<---->	W2
P2 P3 P4 P5		
P6 P6' P6"	<---->	W6
P7 P8 P9 P10 P11 P12		
P13 P13' P13"	<---->	W13
P14 P15 . . .		

Here, I represent word W1 as evoking perceptions P1, P1', and P1", word W2 evoking perceptions P1', P16, and P17, and so on. This sort of partial overlapping and intersecting is typical of word meanings. ("Dog" links through one sort of perception back out to the word "hound" and through another sort of (sexist) perception back out to "broad." Each of these words evokes other perceptions, which in turn have yet other words associated with them.) Conversely, perception P1' evokes both word W1 and word W2.

Now suppose I want to communicate perception P1' to someone. I can pick either word W1 or word W2 (in our simplified universe here). My choice might be determined in part by extraneous perceptions associated with one word but not the other. As Penni and Avery have suggested, my choice is constrained by considerations that apply to the words qua words with respect to other words already spoken or planned to be spoken.

Among these considerations are the operator status of the word, the other words it requires as arguments. Bill argues that these factors are determined by the perceptions (meanings) associated with the words: You can't have a "jump" perception without a perception of a thing jumping. Against this is the proposition that there are only certain kinds of perceptions that you can talk about, namely, those provided for in your language and culture. (You can talk about more inchoate perceptions, but only with difficulty, and with far less assurance of being understood.) Also among these considerations are the reductions available when the word enters on its arguments. This is determined largely by convention. Some redundancy allowing for reduction is determined

There are other ways of describing language that see S, NP, etc. and as byproducts of the simpler elements and relations that they do countenance. However, these other systems are like the familiar phrase-structure-based systems in that they, too postulate ancillary metagrammatical elements and relations as somehow basic.

Harris's aim was a "least grammar", since any additional objects and relations can only obscure the objects and relations by which language "carries" information.

Bruce bn@bbn.com

Date: Fri Nov 06, 1992 5:57 pm PST
Subject: Re: Jeff's Loose control system

[From Jeff Hunter (921106-A)] Re: Bill Powers (921105.1430)

Great. I was afraid when I posted that you would either agree, or disagree with the post. Since you've done both I can be happy :-)

I'll try and state my points again more carefully. If I do it right you should reply "yeah. that's trivially obvious".

First your comments on the "wavers" program. (Which by-the-way are pretty perceptive.)

> I've got your program to run, and think I understand it now. By the
> way, Turbo C doesn't like long integers as indexes to arrays

 Ta. That bin stuff didn't need to be "long" anyways.

> The environmental part of the loop is a simple integrator.

Yup. It's just one-dimensional brownian motion. Fortuitously Hans Blom is also into the topic.

The default change-per-tick is +-4 (just to be greater than 2).

> The perceptual function and comparator are in "if (cont_tot > 0)"...

Yup. Reference is zero. Linear input function. No delay.

The comparator is actually *trinary* since it can produce <0, >0, or =0. This point is obscured by the code, and it makes an analysis of the beast a little tricky.

> The steady-state loop gain of these control systems is nominally infinite.

Great. Glad you agree.

> ... The reason is the integration in the environmental part of the loop.

Not exactly. The reason is that the long term average of the disturbances is zero. If I had biased the random numbers so that there were 3 increases for each decrease then the control would disappear.

This is the main point of "wavers". The effective loop gain of the control can be varied wildly merely by applying a different "pattern of disturbances". For example in the "bias" case the disturbances were +5,-3 rather than +4,-4.

> ... It is
> impossible to estimate the effective loop gain except for the biased
> case, where a constant of 1 was added on each iteration. Over 10,000
> iterations, this should have added 10,000 to the value of bias_tot
> (and a lot more to the squared error); instead, bias_tot changed by
> only 104, or about 1% of the expected value. This yields an effective
> loop gain of about 100.

Forgive me if this point (effective loop gain depends on the disturbances) is obvious. However a lot of discussion on the group appears to ignore this.

I think that the one of the points of the "leaning on the environment" posts was that a perfectly valid control system in one setup could fail spectacularly if the disturbances changed in pattern even if they stayed the same in magnitude.

> If you intended to portray a loose (low-loop-gain) control system,
> you didn't succeed.

I actually wanted to emphasise the difference between moment-by-moment gain (in which "cont" is totally swamped by almost every disturbance) and the long-term gain (in which "cont" has infinite gain).

... Jeff

De apibus semper dubitandum est - Winni Ille Pu

Date: Sat Nov 07, 1992 7:52 am PST
Subject: Re: dear Rick

In reply to: "From Rick Marken (921106.0800)"

>Thanks. You made my morning.

>You are most wise and most welcome to the list.

>How about a quick autobio and/or a comment about what got you onto csg-l?

O.K.

I am a degree candidate in sociology at Texas A&M University writing dissertation entitled "Knowledge, Science, Technology, and Society: a Theoretical Model of Self-Organization." Basically, I am arguing against the prevailing view of functionally differentiated system structures that each act as "externalities" for the other. I am trying to develop a hierarchical model of a self-organized system of resource transfers that utilizes the same unit of analysis for both ecosystems and social systems: the resource transfer group. I borrow this concept from Lee Freese at Washington State University.

In building the model, I am interested in three basic sources of hierarchical control: 1) interpersonal dominance/subordinance; 2) asymmetrical dependencies and interdependencies in "chains" of resource transfers between groups; 3) aggregate control features that may or may not be a perceptual artifact of the mathematical model.

I heard about the csg-1 group via an announcement in the ASA newsletter. I belong to a number of other groups, and decided to take a look. I am interested in the ideas, but very unfamiliar with the area. I appreciate the beginner's bibliography that is posted, and am in the process of locating some of the materials contained therein.

I am as yet undecided how the concept of perceptual control may fit in with my thesis. Thus, I am primarily observing the group discussion at this point.

Thanks for welcoming me to the group. I've enjoyed the discussion so far, even though some of it is Greek to me!

Regards, M. Lee

Date: Sat Nov 07, 1992 8:34 am PST
Subject: Knowledge base; hierarchy & language; address

[From Bill Powers (921107.0800)] Dick Robertson (921105) --

Hardly any of your post got through. When you don't use hard carriage returns at the end of each line, everything gets cut off after the first line until the next hard carriage return. I use my PC to generate text and "print" it to a file, so that the carriage returns are always there (except when I forget and send the original). I hate the editor that the mainframe uses; it's so clunky and non-intuitive and difficult to use.

Penni Sibun(921105.1300) -- Avery Andrews (921106.1348) --

I have some questions about the "knowledge base." I gather that what you mean by a knowledge base is a set of statements describing perceptions, rather than the perceptions themselves. I can see how it would be possible to mark a list of statements as "said" or "not said," but how would you mark the perceptions to which they refer?

In a lot of the proposals offered here by linguists, I get a sense of something missing. PCT itself is about how a brain works (whether we've got it right or not). So I'm always asking "how would the brain do that?" I can understand that people can make lists of things and check them off, but somehow I don't see that as the way language is put together inside the brain -- it's more the way a programmer would do it, or someone using a pencil and paper. I think of a model as showing HOW something is done, even if only in principle; a lot of propositions by linguists seem to me to be concerned with WHAT is done, or proposing ways to use the brain's facilities to achieve an equivalent effect, as in checking off lists of statements.

Avery:

The discrete nature of symbol-handling isn't a problem for control systems. In Rick's spreadsheet analysis, one level of control is logical; the perceptual signal, reference signal, and error signal are all either TRUE or FALSE. The input function computes a logical relationship between continuous variables: $A > B$, and so on (several different control systems working at the same time). If the reference signal for this perception is

TRUE and the perceptual signal is FALSE then the error signal is TRUE. The output function converts a TRUE error signal into a setting (or an increment -- I forget) of a reference signal for a layer of continuous control systems just below; all the control operations in the lower levels of the model are continuous. The higher systems alter the reference levels for the continuous variables A and B until the condition $A > B$ is perceived as TRUE. At that point the perceptual signal equals the reference signal and the error signal becomes FALSE. With the error signal FALSE, the output ceases to integrate and holds its value, so the lower-level reference signals become constant. The lower systems then maintain A and B in the states that correspond to $A > B$ (Rick, do I have this right?). You could let the output decay with time, so eventually the error would become TRUE again. Note that many logical control systems are working at once, so there's no simple relationship between a higher system's output and the net reference signal for a lower system.

You could have several layers of discrete control. Suppose you had a category level that controlled for the perception A OR B OR C. The presence of any of the elements A, B, or C would produce a perceptual signal indicating presence of the category whose elements are A, B, and C. The category control system would have on-off signals in it, and the output would alter reference signals for lower levels (by some scanning or search method?) until at least one member of the category was being controlled by the lower-level continuous system in a sufficiently non-zero state. As a result, a signal indicating presence of that category would exist and be sent to higher systems. This would happen in response to a reference signal, also on-off, saying "provide me a perception of this category" (by using lower systems to create an instance of it). Categories are intersected simply by sending reference signals to more than one category control system at the same time.

If the higher system were a logical control system like the highest level in Rick's model, some logical proposition involving many individual categories would be perceived. The control systems involved would be controlling for the presence of several categories which are variables in a logical proposition. The reference signal could be TRUE or FALSE. The error would be converted to an output that varies reference signals for categories; if one set of categories won't produce the required truth-value of the proposition, the reference signal outputs to the category level are changed from TRUE to FALSE and vice versa until some set of categories is found that can be controlled so as to produce the required proposition with the required value.

This illustrates the basic hierarchical principle in HPCT, which is that one level is concerned ONLY about perceptions of that type, not with the details of how they are derived or what they mean at lower levels. A category control system doesn't care which elements are present, as long as the presence of something in the category is sensed. The logical level doesn't care which categories are involved, as long as they make the logical form TRUE or FALSE as specified by the reference signal.

At the logical level, logical forms are perceived through general functions with places in them for variables. It doesn't matter where the values come from; in the proposition $A \rightarrow B$ (A implies B), it doesn't matter what A and B represent at a lower level. All that matters, if the reference signal is set to TRUE, is that B is not FALSE when A is TRUE. If two categories can be found to plug in where A and B go, and their values turn out to be anything but A present and B absent, the perceived value of the proposition is TRUE and the error is zero. To make the value TRUE, the output of the control system must produce (via category control systems) two categories in the set accepted by that perceptual function that can be controlled in, or are already in, the required states.

When we experience this hierarchy, we see all the levels squashed together in a single space. So it seems to us that it DOES matter what lower-level variables are involved in a

higher-level perception. That is partly because we can perceive as from the lower-level point of view at the same time as from the higher-level point of view. It is also partly because there is more than one higher-level control system involved; lower-level perceptions must be brought to states that are perceived, in parallel, in many different ways, for each of which there is a preferred level of the perception. So a sentence or phrase or conversation, which is a series of words at the lower levels, is perceived both in terms of higher-level meanings and in terms of higher-level forms of speech or writing for which we have preferences.

The hierarchical kind of analysis, I think, can help us separate the different levels that are important in language (as in all else). If we just take what we hear as it comes to us, all these levels are superposed, and it's difficult to see the differences between them. Needless to say, these levels are not those of standard linguistics, such as phoneme, morpheme, word, sentence, and so on -- although there may be some close parallels. The levels are much more general than that: intensity, sensation, configuration, transition, event, relationship, category In language they relate to the features of speech or writing. In other pursuits they relate to other aspects of experience.

I don't mean to suggest that at the "logic" level, the operations are all Boolean, or that at the category level they are just an OR. These are just illustrations; we don't know what really goes on at these levels. The "logical" rules of language are probably not Boolean. I'm just trying to get the principle across, that each level is concerned only with variables of its own type, and that discrete control systems are not hard to understand in general.

Eileen Prince --

Ed Ford failed to say that yearly membership in the CSG for students is \$5. That's one reason the fee for working members is \$45 -- there's a subsidy.

Also, in the address for the Control Systems Group, please be sure to say Ridge PLACE, not Ridge ROAD. And be sure to add the County Road 510 (CR 510), because there IS a Ridge ROAD in Durango West, about 25 miles from here and on the other side of Durango.

Ed, please check your address materials.

Best to all, Bill P.

Date: Sat Nov 07, 1992 1:39 pm PST
Subject: Politics;distributed PCT; Language'

[From Bill Powers (921107.1100)] Rick Marken (921106) --

You have hit a productive note on the political scene. I have wanted for a long time to start a nonpolitical discussion of the systems concepts (etc.) of right versus left, conservative versus liberal, Republican versus Democratic (to the extent that there is a "versus") belief systems. Maybe now that the election is over we can try this on without being accused of using the net for political purposes.

While either side has some problems, I've been intrigued by the simultaneous attraction and repulsion I feel for the conservative view. Obviously a control theorist likes the

idea of independence from external control, reliance on self, responsibility, and non-interference with the freedom of others. These ideals are sometimes beautifully expressed by conservatives and I feel sympathy with them.

On the other hand, conservatives express these ideals, in practical situations, in a way that sounds and often is heartless and selfish. It is true, in PCT terms, that if you help another person to overcome difficulties, that other person will relax and let you do the work. If this is followed as a general policy, people will not learn to fend for themselves. On the other hand, when you have a person in need of help, what should you do? It often seems that conservatives say "Well, it's too bad for that person, but if we made a habit of helping everyone who needs it we would soon be overwhelmed with applicants." What the conservative seems to fear is that helping people in general will end up costing the helper more than the helper is willing to give. A very strong reference level seems to be not to risk losing what you have, even when this means allowing others to suffer.

This is surely a conflict, because conservatives are not usually sociopaths; many of them profess to be sincere Christians. It reflects a theory of behavior, too: the idea that incentives and punishments will change people's behavior. There seems to be a belief that people do not have goals and aspirations of their own -- not the common people, anyway -- and that without some external force they would not try to do anything.

Not being particularly savvy about politics, I'll leave it there as a seed for discussion. There is lots more to say in characterizing the perceptions and reference levels of conservatives, and the same job needs to be done for liberals and the rest. I think that all political points of view represent simply the stronger side of a conflict; every policy has its drawbacks, and no political viewpoint presents a clear picture. And of course, none of them contains any explicit recognition of the properties of living control systems. I would be very interested in everyone's thoughts about this -- not whether one system is better than the other, but just what the perceptions and reference levels seem to be. Maybe if we can get a good look at that subject, we will see how to form a "realist" party.

MARTIN Taylor (921106.1045) --

>...not only is the model not invalidated, it is enhanced by
>distribution. Indeed, I think distribution is the only way that a
>full-blooded PCT system could work in a real, as opposed to a
>modelled world.

I agree with you in the sense that adaptation, learning, self-reorganizing couldn't be handled with the separate-function model. A change in any one function would have to be accompanied by precisely-calibrated changes in all the others in the same cluster. In the real system the nuclei or clusters of functions are, I'm sure, full of interactions; the reorganizations that go on affect many systems at the same time, not just one.

I'm glad you agree about the current connectionist models. I think that the general approach has in it some extremely important ideas and methods. But this approach has to be incorporated in an overall HPCT model, in which there are many more kinds of computations involved at different levels. And connectionists have to stop trying to turn inputs directly into outputs; they're missing the whole point of the control organization because of a basically S-R orientation left over from older theories. In the brain, input processes are quite distinct from output processes at the lower levels (up to the interior layers of the cerebral cortex) and one can even find groups of comparators. Inside the cortex it's not so clear any more, but at least in the visual systems there are well-defined stages of processing that are clearly input processes.

I would like to see a compromise between the general approach now being taken and my simplistic modeling. Inside a perceptual function, or perhaps even a layer of perceptual functions, the connectionist approach could be used to explain how input functions become organized, and also output functions. What I visualize is a smallish input network with multiple inputs and multiple outputs. The multiple outputs carry perceptual signals upward to the next level and laterally to comparators associated with an output network which has its own, and perhaps quite different, principles of self- organization. For the first four or five levels, at least, this picture would be consistent with neuroanatomy, in which input and output processes are physically distinct and distant from each other and are connected only by signals. A form recognizer doesn't need to produce behavioral outputs. All it needs to do, with respect to our modeling requirements, is produce a set of perceptual signals.

Bruce Nevin (921102.0738) --

Nice to see you back in stride again. Something seems to have taken more shape in your investigations of PCT and linguistics. It's beginning to sound more coherent to a non-linguist. I think that some of our pseudo-disagreements simply went away during your preoccupation with other matters (as much in me as in you).

You might recall that last Spring (or sometime) I offered a proposal about how communications were formed by a speaker. This was in the context of Harris' operator grammar. I was looking then, and am still looking, for the reasons behind the rules for operators and arguments, reasons other than just "that's how they say it." In the course of this I got the idea that AS we construct sentences or strings of words, we are concurrently perceiving the meanings that they evoke in ourselves, and progressively making sure that the sentence to date does not create any important meaning errors. Your sets of "do say" and "don't say" sentences seem to illustrate this rather clearly:

>One might suppose this is because "donate to" simply is not used as
>frequently as "give to." It is true that the dative reduction is
>more or less peculiar for many other verbs in the "give" set. Call
>them the "donate" subset, including:

? John administered [to] the fund the money
? John allocated [to] the fund the money
? John apportioned [to] the fund the money
? John communicated [to] the fund the money
? John consigned [to] the fund the money
? John contributed [to] the fund the money
? John conveyed [to] the fund the money
? John dealt [to] the fund the money
? John delivered [to] the fund the money
? John dispensed [to] the fund the money
? John distributed [to] the fund the money
? John entrusted [to] the fund the money
? John purveyed [to] the fund the money
? John relinquished [to] the fund the money
? John rendered [to] the fund the money
? John surrendered [to] the fund the money
? John transferred [to] the fund the money
? John transmitted [to] the fund the money
? John vouchsafed [to] the fund the money

>But a considerable number of other verbs in the "give" set are
>acceptable with the transposed argument order and zeroed "to".
>Call them the "give" subset:

John accorded [to] the fund the money
John allotted [to] the fund the money
John assigned [to] the fund the money
John awarded [to] the fund the money
John conceded [to] the fund the money
John furnished [to] the fund the money
John granted [to] the fund the money
John handed [to] the fund the money
John left [to] the fund the money
John lent [to] the fund the money
John offered [to] the fund the money
John paid [to] the fund the money
John presented [to] the fund the money
John proffered [to] the fund the money
John provided [to] the fund the money
John sent [to] the fund the money
John supplied [to] the fund the money John yielded [to] the fund the money

>I don't think a "frequency of occurrence" argument holds up.

I don't either. But looking over these two sets, I was reminded of my earlier proposal.
If you read slowly, attending to meaning errors as you go, look what happens:

From the first set:

? John administered the fund
? John allocated the fund
? John apportioned the fund
? John communicated the fund
? John consigned the fund
? John contributed the fund
? John conveyed the fund
? John dealt the fund
? John delivered the fund
? John dispensed the fund
? John distributed the fund
? John entrusted the fund
? John purveyed the fund
? John relinquished the fund
? John rendered the fund
? John surrendered the fund
? John transferred the fund
? John transmitted the fund
? John vouchsafed the fund

From the second set:

John accorded the fund
John allotted the fund
John assigned the fund
John awarded the fund

John conceded the fund
John furnished the fund
John granted the fund
John handed the fund
John left the fund
John lent the fund
John offered the fund
John paid the fund
John presented the fund
John proffered the fund
John provided the fund
John sent the fund
John supplied the fund
John yielded the fund

Both of these sets of statements create a large meaning error when they get to the stage shown. The meaning in each case, but more so in the first set, is clear but WRONG, because the intended meaning is not that the verb is being done to the fund. At this point in the sentence, however, it strongly seems to apply to the fund. In the second set there is more of an error of usage than in the first set, as you imply. But both sets create a positively wrong meaning before they are finished.

On the other hand, if you stick "the money to" in the middle, there is no meaning error at all as the sentence progresses; only what error is present because of non-completion. You could put a period after "the money" and all the sentences would still have correct meanings, although some added meanings (the destination of the money) remain to be described.

If you happen to start any of these sentences in this way, you can't terminate the sentences where shown without conveying and experiencing the wrong meaning. You MUST tack on "the money" to make your meaning clear, but even when you do, the echo of that midstream meaning error is still present. The sentence feels awkward with "the money" on the end. The only real cure is to rearrange the sentence so it does not create meaning errors at any point during its utterance.

This implies, of course, that meaning (plus convention) dictates the order in which we connect words together, and does so during the construction of utterances.

RE: the diagram

>This diagram makes the point about arriving at outputs by continuous
>control of perception rather than by top-down realization of some sort
>of archetype. The actual flow of control is probably quite different.

No doubt. As you say, I am glossing over some very complex processes in this diagram. You could unpack all the functions a lot, if you knew what they need to be (I don't). Associative memory must play a large part, which can't be shown in a single-loop diagram.

Jeff Hunter (921106-A)--

> The comparator is actually *trinary* since it can produce ><0,
>0, or =0. This point is obscured by the code, and it makes an
>analysis of the beast a little tricky.

OK, I missed that. No adding or subtracting if the error is exactly

zero.

```
>> ... The reason is the integration in the environmental part of
>> the loop.
> Not exactly. The reason is that the long term average of the
>disturbances is zero. If I had biased the random numbers so that
>there were 3 increases for each decrease then the control would
>disappear.
```

Well, I just tried it. I did this:

```
blip = random(0x200) & 0x180; /*NOTE - USING 2 BITS*/
if (blip == 0x100) {
    d = increment;
    pos++;
    if (last) high++; else up++;
} else {
    d = -increment;
    neg++;
    if (last) down++; else low++;
}
last = blip;
```

My random function generates integers between 0 and the argument. So this makes three decrements occur for every increment, on the average.

I then varied the "increment" (the disturbance) and printed out just the final value of the controlled (or free) variables. I did this for both the original case and the case with the offset. Here is the result:

ORIGINAL CASE: EQUAL INCREMENTS AND DECREMENTS

Increment size	free_tot	cont_tot	bias_tot
1	-16	0	0
2	-32	0	2
3	-48	0	2
4	-64	-8	0
5	-80	-18	0
6	-96	-28	-10

3 DECREMENTS PER INCREMENT

1	-5000	0	0
2	-10000	0	2
3	-15000	-12	2
4	-20000	-376	-8
5	-25000	-5008	-34
6	-30000	-10004	-564

In both cases, where the final values of cont_tot and bias_tot begin to change, the reason is that control has been lost altogether and the values of these variables are increasing without limit. The longer you run the bigger the error gets, so there is no control at all. This begins to occur at a value of the increment where the average

disturbance added on each iteration exceeds the amount of correction that the control system can add to the controlled variable.

You will notice that when the average disturbance is smaller than the maximum that the control system can resist, control is perfect. Your way of representing the data conceals this fact, because the numbers you get represent only the entire data set, not its behavior through time. You should have been looking at the controlled variables in detail, not using measures representing the whole run.

I realized at this point that you were misled because of using only a reference level of zero. If effect, you are looking at the noise in the control system and missing the main effect. Here is a run of the second condition above with the reference level set to 10 by statements of the form

```
if (*_tot > 10) ...  
if (*_tot < 10) ...
```

3 DECREMENTS PER INCREMENT

Increment size	free_tot	cont_tot	bias_tot
1	-5000	10	10
2	-10000	10	12
3	-15000	-2	12
4	-20000	-336	2
5	-25000	-5000	-24
6	-30000	-10000	-554

Now it's clear that an RMS fluctuation of less than one unit is negligible relative to the reference level, here 10 units. For the case with 1 unit of random disturbance, the RMS fluctuations are actually about 0.001 units, or one ten thousandth of the reference value. You couldn't see this with a reference level of zero because you had nothing to compare the fluctuations against.

Under this second condition, the average value of the disturbance is $(3 \times \text{increment} - \text{increment})/2$ or simply the increment size. You see that control of `cont_tot` is perfect for an increment size of 2, which is just equal to the maximum output of the control system. Actually statistical variations would make this not quite perfect control, but with integers we can't tell the difference from perfection. For increments of 3 or greater, control is lost; the controlled variable is actually changing continuously, and with continued iterations would become larger without limit. Note that with an increment of 6, the average value is -3; add the +2 which is the maximum possible opposing output, and you get an average of -1 unit of unopposed disturbance per iteration, accounting for the final value of `cont_tot` of -10,000 in 10,000 iterations.

So the difference in control with the disturbance present is not due to a change in loop gain, but to using disturbances too large to resist. The loop gain is zero for average disturbances greater than + or - 2. This shows the pitfalls of estimating apparent loop gain without looking at the details of behavior -- in fact, there was no control at all, but the data averaged over 10,000 iterations concealed that fact.

Best to all, Bill P.

9211B

Date: Sun Nov 08, 1992 11:29 am PST
Subject: internal conflict

from Ed Ford (921108:1220) Tom Bourbon (921106)....

>.....I believe Ed meant that he would like to call down a pox on both
>parties and on both the congress and past presidents. (Is that it, Ed?)..

That's it exactly Tom. I am reminded of the famous line from the movie A Man For All Seasons. Thomas More (Paul Scofield) is resisting Cardinal Wolsey's (Orson Welles) request for help in pressuring Rome to give in to King Henry VIII's desire for an annulment from his first wife by threatening to take or tax church property. Wolsey asks "Explain how you as counselor of England can obstruct these measures for the sake of your own private conscience." I think More's answer is a classic, and the reason we are experiencing the mess we are in. He replies "Well, I think that when statesmen forsake their own private conscience for the sake of their public duties, they lead their country by a short route to chaos."

It seems to me that by forsaking our own systems concepts, that which we've built over a life time of thinking and experiencing, for the sake of what we perceive as social demands or social approval, as More states, we lead our country as well as ourselves "by a short route to chaos." I would add that we may also be placing ourselves in conflict if we are making choices that are in conflict with our systems of beliefs which we hold sacred. We may have two areas of importance within our systems concepts that are incompatible. On the first side, the approval of the voters and the power that comes with the office and thus the decisions we make that help us maintain this status quo; on the other side, other systems of beliefs that we've formed and held in high respect over the years that are now in conflict with the first side.

Closed Loop was mailed out Saturday to paid up members of the CSG.

Best, Ed.

Date: Mon Nov 09, 1992 9:27 am PST
Subject: Elections

[From Hank Folson (921108)]

I am fascinated by the opportunity to compare & contrast S-R (Stimulus- Response) and PCT interpretations of the elections. One example: Ross Perot spends \$60,000,000 on his campaign and gets 19% of the vote. S-R interpretation (taken from the news media): Perot's campaign/money/ads "stimulated" one in 5 voters to "respond" with Perot votes. PCT interpretation: 1 in 5 voters had large error signals with Republicans/Bush and Democrats/Clinton. Perot's personality/policies/ads matched their reference levels better so they voted for Perot.

How will the political parties, the media, the large corporations interpret the meaning of this vote? It seems to me that if they accept the S-R view, they will discount the meaning of the Perot vote. They will feel that the problem is simply one person, Perot, "stimulating" those voters. If they don't like what happened, they will concentrate on doing something to counteract Perot's future "stimulations". They will consider those who voted only in terms of their responses to future stimulations. If they take the PCT

view that these voters were controlling to reduce error signals, they will look at Perot only in terms of how his campaign was perceived by these voters. The actions of the politicians, media and corporations will be different depending on which philosophy they follow, and their success in achieving what they want will be different, too!

Do any other examples of different S-R vs. PCT interpretations and consequences come to mind?

There was some prediction of our political future in earlier posts. How can we begin to make any predictions unless we have some awareness of the settings of the reference levels of the higher levels of our politicians?

Hank Folson (non-voting Resident Alien)

Date: Mon Nov 09, 1992 9:28 am PST
Subject: language, discreteness

[Avery Andrews (921108.1244)] (Bill Powers (921107.0800))

> I have some questions about the "knowledge base." I gather that what
> you mean by a knowledge base is a set of statements describing
> perceptions, rather than the perceptions themselves. I can see how it
> would be possible to mark a list of statements as "said" or "not
> said," but how would you mark the perceptions to which they refer?

The perceptions being described are typically long-gone in these situations, so I don't see that there's any sense in which they can be marked at all. It's my impression, though I don't know the area, is that there is a certain amount of evidence that memory really is symbolic to a considerable degree, and so, when one is talking on the basis of what one remembers, it is quite appropriate to mark statements as 'said' or 'not said', at least as a first approximation.

As for modelling, what you are sensing missing is any idea of how linguistic tasks could be done with brain-like-as-we-know-it hardware. The tricky bit, I think, is how to deal with the fact that we seem to be able to build up networks in which arbitrarily large numbers of individuals are classified and related with a finite number of properties and relationships. E.g. we can learn of any number of women that they are daughters of, say, Adeline, any number of men that they are (current or former) husbands of these women, etc. & people can learn this kind of stuff very quickly--much faster than a connectionist network could be trained.

I don't think anyone has a clue as to how to do this with neurologically realistic hardware (wetware?), just like nobody had any idea in the nineteenth century about how to reduce chemistry to physics, though people seemed to have assumed that it had to be able to be done somehow. Maybe this means that PCT should ignore language entirely until this problem is solved, but I think that the idea of closing loops through the environment has enough power to be worth trying to apply to language even without neurologically concrete modelling.

And, it seems to me that trying to build models with the emphasis on interaction & minimal representation might turn up some constraints that would make a neurologically realistic model easier to attain. At any rate, it's the best we can do at the moment.

The point of my little piece on discreteness is not that it is not necessarily problematic for PCT, but that a different collection of things might be important in a regime where disturbances are highly limited (e.g. once something has been said, nothing can make it unsaid), and where reference levels can be attained *exactly*. In particular, it is not at all clear to me that the usual equations have much significance--what makes equations interesting is the possibility of solving them non-trivially, but I don't see how to do this in this sort of domain: the perception that cancels the error signal is to hear somebody saying what is to be said, & that's a lot less exciting than the notion of the equilibrium state of a feedback system, at least to me (since the solution is a trivial restatement of the problem).

Avery.Andrews@anu.edu.au

Date: Mon Nov 09, 1992 10:08 am PST
Subject: Question on Demo. Promoting PCT.

[From Dag Forssell (921108)]

The phenomenon of control:

While on the subject of promoting PCT, I would like to ask out loud: Have any of the people who requested and received my tape and demonstration packet looked at it? I have not heard a single pipsqueak. Did anyone read Ed's book? Do you understand the phenomenon of control any different than before?

Renewed efforts: Christine and I feel as if we are on a higher plateau of progress now. We are renewing our marketing efforts with a clearer sense of what we offer. This next piece is part of that.

Here is a draft for a promotional paper on Purposeful Leadership and PCT. Any comment will be appreciated. I am anxious to be correct in what I say, and as gentle as possible while stating the case forcefully for PCT.

UNDERSTANDING A THEORY-BASED PROGRAM

CONTENT:	Page
Theory, REALITY, training, scientific revolutions and PCT	1
Who will pay attention?	6
Who can understand?	6
PCT applied in values, methodology and role play	7
Some of what you will learn	9
Relevance to work life	9
Relevance to personal life	9

THEORY, REALITY, TRAINING, SCIENTIFIC REVOLUTIONS AND PCT

Bertrand Russell, the famous British philosopher, is reputed to have said:

There is nothing as practical as a good theory.

In the engineering and physical sciences, this is obvious. Engineers and physical scientists recognize that a tested theory allows for the prediction of performance long

before actual experiment or production. Good theories have allowed us to travel, communicate, understand and produce better than ever before.

The power and practicality of a good theory is awesome. It also provides clarity.

When it comes to the important area of human affairs, this is not quite so obvious. Many theories have been offered over the years, attempting to explain human action, but none have measured up. This is why many psychologists say that their theories and practices have nothing to do with each other.

Since the results have been poor, scientists have settled for descriptions and statistical correlations between described phenomena. Explanations which yield predictions with 99+% experimental confirmation have not been possible and are not expected. Levels of explanation and proof required for publication are orders of magnitude lower than in the engineering sciences.

Some say this is because humans are too complicated. Could it be that the concepts have been inadequate? As we shall see, the history of science is full of phenomena that were considered mysterious and too complicated until an appropriate concept was offered.

Let us distinguish between the objective physical REALITY of the world, and any one person's subjective interpretation of it: personal reality.

REALITY: What actually exists (but can never be known).

reality: What a person sees, hears, touches and smells.

Personal, subjective reality is like a map of REALITY. The personal reality is all a person can know. I think of reality as a heads-up display in a fighter pilot's helmet, except it is internal to our own head. It is only a display. REALITY is beyond our grasp, but we spend a lifetime trying to understand it. As we do, the subjective map of reality is formed in the brain and constitutes a personal world.

To a person who looks at a seamless world in living color and stereophonic sound, this distinction may seem silly. Obviously, what we see, hear, touch and smell is REAL, isn't it? Careful thought tells you that it cannot be. Your senses may be well calibrated and yield a very good map of the physical world, but a moments reflection tells you that it can only be a map. All you know is an accumulation of interpretations of signals from nerve endings. Today we know that there are many things in existence which we as humans can not sense at all. Elephant-talk below 18 hertz, echo-location beyond 20,000 hertz, infrared light, X-rays; all are part of a long list of phenomena we cannot sense. Clearly, our sense of the world is just that, sensed perception, a display.

Consider how a person grows up and develops a personal understanding of how the world "works" from personal experience, supplemented by experiences told by others, seen on TV and read in books. (PCT suggests how). Each person develops a unique framework of ideas, that personal reality which governs that person's life. These personal ideas will vary. The quality of the personal map reality will vary widely from person to person.

A personal map of high quality, a good reality, allows us to navigate successfully in that elusive REALITY the map portrays. A good theory is a good reality; a map of a specific aspect of REALITY, expressed as description, explanation and perhaps a set of clear rules for prediction that follow from the explanation. Many people can easily share the same good piece of reality. This is why the power of a good theory is awesome.

Consider programs for leadership/management/teamwork training. They are based on "what seems to work" - anecdotal experiences by the authors, experiences from many sources and conclusions drawn by the authors. These programs of course each portray a reality, but in the absence of proven theories in the area of human affairs they cannot offer a universal framework of explanation. The focus and quality varies.

Companies spend millions of dollars on training. In many cases, people have fun and like the training, but four or five months later, little has changed in the workplace. I believe a large reason for this is that the training is situational / anecdotal and focuses on "what/how to DO." Each participant is left to integrate the many disparate lessons of the training experience into the framework of their personal reality or understanding, such as it is.

People ask: Show me what to DO (cause) so I will get results (effect). Our program shows clearly why this is fallacious, but that does not change the fact that this is what many people expect until they know better.

Instructions on "what/how to DO" are valid only in a given set of circumstances. Typically a training scenario is carefully selected and told with drama and humor by a speaker. You are told what was done and what the results were. You imagine that the same thing will happen if you do the same thing. You feel euphoric as you imagine success.

A large part of the "what/how to DO" training does not really apply in individual cases because the world is full of varying conditions and changing disturbances. The lessons become irrelevant and are soon forgotten. Euphoria fades away.

If much training is ineffective, how can we suggest that Purposeful Leadership is of lasting value?

The major strength of the Purposeful Leadership program is that it explains and applies a new theory called Perceptual Control Theory (PCT). PCT recognizes and explains the phenomenon of control. It explains why and how people do what they do. PCT is based on neurology and clear, detailed and tested engineering concepts. PCT requires and offers scientific rigor with explanation and prediction.

PCT is an engineering science of psychology that is easy to understand for an engineer used to explaining/predicting/testing scientific thinking.

Once the phenomenon of control is observed and the detailed explanation understood, it will be seen that control is the fundamental organizing principle of life. Control is pervasive and can be seen operating at a microscopic level as well as at the macro level of human activity.

PCT explains a wide variety of phenomena of everyday experience.

If this new theory is so much better, why is it not widely known already?

Thomas S. Kuhn, a leading scientific philosopher, professor at MIT and author of: The Structure of Scientific Revolutions, provides an answer. From the book cover:

....Thomas S. Kuhn wastes little time on demolishing the logical empiricist view of science as an objective progression toward the truth. Instead, he erects from the ground up a structure in which science is seen to be heavily influenced by nonrational procedures...Science is not the steady, cumulative

acquisition of knowledge that is portrayed in our textbooks.
Rather it is a series of peaceful interludes punctuated by
intellectually violent revolutions ...in each of which one
conceptual world view is replaced by another...

Nicholas Wade, Science

Thomas Kuhn introduced the term "paradigm" in this book and suggests that scientists schooled in a certain set of views adopt them as their personal "paradigms" (or reality) and then view the world through these paradigms - as if they were eyeglasses filtering information. The word paradigm means pattern. It is used to signify how we interpret a phenomenon; how we explain the world to ourselves.

In Kuhn's view, everyone is a scientist, and every world view (or reality) might be called a personal science. Everyone has some framework of ideas of how the world "works" and views the world through those personal paradigms.

There are many levels of scientific rigor and interpretations of the word theory.

To some, "theory" is a "four letter word" signifying head-in-the-clouds irrelevant hokus-pokus. All theory is incomprehensible and irrelevant. From this perspective: You know what they say about engineers.

To some, "theory" means a description of a phenomenon. It is hard to test descriptions of phenomena in a rigorous way. Words are defined by other words.

To some, "theory" means a logical structure built on well defined hypotheses, which stand apart from the real world. This kind of theory is subject to logical proof.

To some, "theory" means precise description, explanation, prediction and test. Testing is done in the "real world" of hardware and is subject to small errors due to things like manufacture and measurements.

A person can certainly be aware of different meanings of the word theory and use it differently in different contexts.

To me, theory is the information / understanding / expectation a person uses to guide activity, no matter what the person calls it or how clear / valid it is as a map of that REALITY. In this view, everyone is full of theories and nothing but. But the quality varies.

When a new theory (normally an extension of existing theory) is offered to an engineering scientist, the course is clear. Study the theory based on existing paradigms (which have proven 99.999+% dependable and usually allow excellent comprehension). Predict something based on the new theory, perform the experiment and accept or reject the theory expecting 99.999+% performance.

When a radically new theory is offered to a scientist, understanding can be difficult. If the old and the new concepts are incompatible, it may literally be impossible to see the new theory through the eyes of the old paradigm.

To illustrate, let us go back to a well known episode of a revolution between a well established descriptive science of astronomy and a new, incompatible, engineering science of astronomy.

In the beginning of the 17th century, Ptolemy's model (AD 140) of an earth-centered universe was still fully accepted. This model was complete with crystal spheres to carry the heavenly bodies and ample room for heaven and hell beyond the eighth sphere, the firmament. Learned scholars who had studied and accepted this model and fitted it into their reality, were not open to a new and different explanation.

In 1543, on his deathbed, Copernicus published a model of a sun-centered universe. (Kepler published additions in 1609-18). In 1610 and more explicitly in 1632, Galileo supported Copernicus' theory, based on the first ever observations by telescope and the discovery of Jupiter's moons. He had trouble being published in the scientific journals of the day (church bulletins?), due to the requirement that his manuscript be accepted by a group of his peers.

(It is of passing interest that on October 31, 1992, Pope John Paul II acknowledged the church's error in this matter).

We understand today that his ideas were incomprehensible to the (church) scholars of the day, due to what they already understood as their personal paradigm. Of course Galileo's writings were irrelevant to their science as they understood it and not a constructive and welcome contribution to the state of the art.

So Galileo self-published. Despite the persecution he personally suffered as a result, astronomy started over as an engineering science and made rapid progress.

Isaac Newton's *Philosophiae Naturalis Principia Mathematica*, (1687) completed the revolution, but was resisted in the same way. According to Kuhn, it took fifty years for Newton's work to be fully accepted.

This well known revolution is not an isolated case. Kuhn describes many upheavals in several disciplines. The opportunity for a revolution arises when a current paradigm fails to solve problems and competing paradigms are offered to provide better explanations. A struggle of many decades typically takes place. Established scientists continue the development, application and defense of the existing paradigm as usual while outsiders and early converts champion a new one.

A reading of Kuhn's seminal book makes it clear that there is lots of room in our sciences of today for coming revolutions.

New information - on any subject - is always filtered by what you already understand. PCT itself explains why this is so. Where a person has existing convictions, conflicting information is either not comprehended or rejected.

A person without convictions on a certain subject is more open to new information.

This is why scientific revolutions typically originate from outside the scientific community which has accepted the present paradigm.

Perceptual Control Theory is a new engineering science of psychology. It offers description, explanation and prediction. Tests show 95-98% correlation, with the remaining 2-5% accounted for by expected imperfection of control: less than infinite loop gain, slow response and sloppy connections in the environment.

PCT offers an opportunity for a transition from a descriptive science of psychology where theory and application are worlds apart to an engineering science of psychology, where theory and application fit like hand in glove.

For contemporary journals of psychology to publish articles on PCT, the phenomenon of control must be understood in detail in a review by peers who have internalized an understanding and strong convictions based on a scientific method limited to the study of cause and effect. An understanding, endorsement and acceptance of PCT is obviously very difficult. PCT itself explains why.

On the other hand, PCT is immediately acceptable -intuitively obvious- to people without such understanding and convictions. It is easy to understand and immediately useful.

Without good theory, every problem must be solved by trial and error.

In engineering school, you learn a theory, then spend time on problem after problem to learn to recognize how the theory applies and get used to using it. What you remember is the theory, not individual solutions.

With Purposeful Leadership and PCT, you learn a theory/structure, then spend time with application after application to recognize how the theory applies and get used to think that way. What you remember is the theory, not individual "what/how to DO" solutions.

The power of a program based on a good theory is awesome.

There is nothing as practical as a good theory based program.

WHO WILL PAY ATTENTION?

PCT explains clearly that the only thing that drives us are our concerns. For this reason, the program, theory and applications must be understood to have relevance to the concerns that are in the forefront of the mind of each participant.

Those participants who have concerns which are addressed by the things PCT promises to explain will want to pay attention. Those who are satisfied with their own understanding, skills and results have no reason to pay attention.

WHO CAN UNDERSTAND?

Because personal concerns determine what is of interest and relevant, the program is structured to provide immediate relevance and usefulness as much as possible, with maximum clarity: (Program details shown elsewhere).

1st Day: Introduction to theory, values and methodology
Effectiveness application

2nd Day: Related concepts
Leadership applications

The first and second day require careful attention, nothing more.

3rd Day: Biology / engineering details.
Structural details.
Computer demonstrations.

Again, attention and interest is the key. A willingness to think in terms of explanation, prediction and test is very helpful.

PCT APPLIED IN VALUES, METHODOLOGY AND ROLE PLAY

The value of Purposeful Leadership is in the theory we teach. To make the theory come alive with relevance to the individual, it must be applied. As the chinese proverb says:

I hear	- I forget.
I see	- I remember.
I do	- I understand.

PCT describes how an individual operates in all circumstances. (We illustrate both cooperation and conflict in active demonstrations). PCT makes no value judgements.

But most people adopt values which constrain their options in working with others.

We specifically define (using PCT) what we mean by:

Supporting individual effectiveness

Respect for the individual

Balance: Appropriate level of direction and involvement
- minimize conflict
- maximize cooperation

Given an acceptance of these leadership values/objectives, we offer a comprehensive methodology which can serve as a guide in a great number of circumstances.

We illustrate the theory, using this methodology, in several scripted vignettes. These are read, broken down to show the use of the methodology and discussed among participants.

The variety of interpretations and applications in these role plays makes it clear that the methodology is NOT a rigid "what/how to DO" set of rules, but only a guide to help you systematically consider various aspects of what the theory teaches you. Situations and disturbances in real life vary so much that the action required to accomplish the same outcome is never the same. This is why it is fallacious to teach "what/how to DO."

Some participants will view the methodology as a comprehensive "what/how to DO" prescription despite our comments to the contrary. (After all that is what people have come to expect from seminars in a cause-effect world). There is no harm in this. Taken as that, the methodology is an excellent "what/how to DO" prescription.

Participants are then invited to plan an act out a sequence of similar situations, given a scenario which is relevant to them, taken from their world of daily work. This scenario is developed ahead of time by a representative of the group and the teacher.

Each role play for a participant provides practice for a work situation, the way a sales role play provides practice for the salesman's work. This makes it personally relevant. The scenario for role plays is tied to the work environment where the training is to be

applied, and provides continuity between role plays, from one to the next. Additional details will come from the common experience and imagination of participants.

This detailed scenario will do triple or quadruple duty. It portrays (for each group) some common process, problem or source of conflict between the fictitious characters such as Joe, Dave and Bill which can :

- 1) be resolved between Joe and Dave in a manager/subordinate or peer to peer conflict role play, focusing on wants, then
- 2) be facilitated by Joe in a team building situation between Dave and Bill, focusing on mutual perceptions, then later
- 3) proudly reported as an accomplishment by the employee Dave or brought up as a concern by the manager Joe in a performance review and finally
- 4) serve as a concern Dave has, addressed by the salesman Joe. (Sales only if a sales application is desirable given a particular audience).

Role play 1 is part of the agenda for Day 1, role plays 2-4 are on the agenda for Day 2.

SOME OF WHAT YOU WILL LEARN

How and why all understanding is individually subjective.
How and why what you choose to want affects what you do.
How and why what you choose to perceive affects what you do.
How and why dissatisfaction is the mother of invention.
What motivation is.
How to develop and support individual effectiveness.
Why two people can look at the same facts and draw different conclusions.
How and why your memories affect what you do.
What it means to respect another person.
How and why stimulus-response is an illusion.
How and why a reward may not motivate but create resentment instead.
When helping a person conflicts with respect for the person.
How to communicate effectively with subordinates.
How to conduct mutually satisfying and productive performance reviews.
What is required to develop and maintain an effective team effort.
How to resolve conflict with another person if at all possible.

RELEVANCE TO WORK LIFE

Conflict resolution: - in a superior / subordinate relationship.
- in a peer - to - peer relationship.
- in a subordinate / superior relationship.
Total Quality Management framework.
Leadership understanding.
Team development.
Performance reviews.
Vision/Mission statements.
Goal setting.
Non-manipulative selling.
Teaching effectiveness, initiative and responsibility.

Develop trust, mutual respect and high morale.

RELEVANCE TO PERSONAL LIFE

Conflict resolution.
Develop loving relationships.
Develop self-sufficient children.
Develop self confidence.

Best to all, Dag

Date: Mon Nov 09, 1992 10:56 am PST
Subject: Language

[From Hank Folson (921109)] (Bill Powers (921107.0800)):

>...linguists, I get a sense of something missing. PCT itself is about how a
>brain works...

If PCT is correct, language is just one of the tools we use for controlling our perceptions. If language suddenly ceased to help us to control our perceptions, we would not continue to use it. I don't think linguists see language in such a cold-hearted way.

What I feel is missing in the linguistic posts is a linking of the linguistic theories and examples to the nature of control systems. For instance a recent post quite logically explained HOW languages can change over time. But there was no mention of WHY in terms of PCT. My guess in PCT terms is that a living control system will not control for zero change in words and their pronunciation unless the changes interfere with controlling actions. Thus language details can drift any which way as long as the changes occur slowly enough that they are not disruptive.

>a lot of propositions by linguists seem to me to be concerned with WHAT is
>done, or proposing ways to use the brain's facilities to achieve an
>equivalent effect, as in checking off lists of statements.

What else can they do? Linguists do not have PCT as a recognized and accepted tool for understanding how the human mind is structured and works. As a result, they have to make do with what they know. It is as if the mind is a "Black Box" that they have no access to. All their understanding must come from observing only the language output and not the machinery that generates the language.

My feeling is that to understand language, rather than starting from the spoken word and working backwards, it will be easier to start from how the human mind functions and work forward using PCT. This will be very difficult for linguists to get into, if linguists have a systems concept of language as an entity that evolves and has a life of its own.

Date: Mon Nov 09, 1992 11:21 am PST
Subject: PCT, Politics;learning

[From Rick Marken (921109.0900)]

While walking with me wife this weekend she challenged me to come up with some surprising conclusion derived from PCT that might grab the interest of a layperson. Here's what I came up with.

"You can't tell what people are doing by watching what they are doing".

It's not often that I say stuff about PCT that grabs her attention -- but she really liked this one. So I offer it to the net for comment.

>Bill Powers (921107.1100) -- Rick Marken (921106) --
> You have hit a productive note on the political scene.

I would like to talk more about the system concepts of the right and left; including what I think are the implicit models of human nature that are a part of these concepts. But there was an article in the Sunday Times that made me realize that there is an implication of PCT for politics (or policy) that we have not really considered yet. I'll say what it is -- then describe the article that motivated it: PCT suggests that it is the GOAL, not the MEANS, that should be important in politics. What we should do is try to articulate the kind of society we want, NOT how we should get there. The latter "depends" on the current "context" of society (ie. disturbances and constraints). So there can't be one "right" way to produce a particular societal result -- just as there cannot be one "right" way to move the handle to keep the cursor on target. The means MUST vary to produce the intended result. Social ideologies are generally descriptions of policies which are supposed to achieve some vague goal -- like a better society. But it is clear from looking at the world that the results of various policies have always been different. Sometimes socialistic practices "work" -- sometimes they fail. Sometimes free market practices "work" -- sometimes they fail. PCT suggests that this is to be expected -- it's not that we cannot know what to do to produce a particular social result (though this is close to being true); more important, continuous contextual changes in society (and the environment) mean that a particular means of producing a particular result CANNOT POSSIBLY be expected to work all the time; the means must vary in order to produce consistent results.

The article that got me thinking about this was in the Business section (which I rarely read). It was about how "history may judge Reaganomics very harshly". I, of course, read and enjoyed the article since it pointed out some of the cruel results of the Reagan era. Now one can try to blame the Reagan disaster (my kids and grandkids will be paying for it for years) on congress or communists or whatnot. But the fact of the matter is that the main agenda of the Reagan era was passed happily by congress -- massive tax cuts, especially for the upper tax brackets. The effect was supposed to be increased revenue due to all that economic growth -- ie. trickle down. This didn't happen (to a disastrous extent) -- thus, a massive debt was born.

What struck me was that the "trickle down" idea is not necessarily "intrinsically" wrong. It might have worked -- maybe it did work some- where at some other time. The problem was that when it became obvious that it was NOT working there was no effort to take any steps that MIGHT remedy the situation. Their was a BELIEF about how the MEANS to an end SHOULD work -- beliefs are just references that are not altered in order to achieve higher level goals.

It may be (as I suggested in an earlier post) that the Reagan/Bush crowd HAD NO higher level goals; tax cuts WERE the highest level for them (they seem to be the highest level for most people). I hope that we can move politics to a level where we are trying to articulate system concepts rather than yelling at each other saying "this is the RIGHT

way to do things". I believe PCT shows that this is just wrong-- there can be no RIGHT way to achieve a particular social goal -- no matter WHAT that goal might be.

That's why I like Clinton -- he seems to understand that it's the vision that matters, not the means (he basically said that in the debates). You have to be willing to try various things and see whether the result is moving you toward or away from the goal. But without a goal, there is no way to judge whether or not you are "doing the right thing". I think the "vision thing" is the whole ball game; the vision thing means "having a reference level for a particular system concept". That's what's important. What you do to achieve that vision MUST BE determined by "context" -- not belief.

Bill Powers to Jeff Hunter --

> So the difference in control with the disturbance present is not due
> to a change in loop gain, but to using disturbances too large to
> resist. The loop gain is zero for average disturbances greater than +
> or - 2. This shows the pitfalls of estimating apparent loop gain
> without looking at the details of behavior -- in fact, there was no
> control at all, but the data averaged over 10,000 iterations
> concealed that fact.

Thanks for doing the work, Bill.

Jeff. Again, I think it's great that you are exploring the control model through simulation. But do try to read some of the literature that I suggested; it can save you from some of these blind alleys.

Best regards Rick

Date: Mon Nov 09, 1992 11:29 am PST
Subject: Primer for modelers: draft

Experimenting with the control paradigm
A primer for computer modelers
DRAFT: William T. Powers

PART I: PROPERTIES OF A SIMPLE STABLE CONTROL SYSTEM

The following is intended to introduce programmers and control-system engineers to the terminology and architecture of control theory as used under the name PCT, or perceptual control theory. For those who have experience with modeling control systems, there are some readjustments to be made, because we will divide the system into functions in a way that is not standard in control engineering. For example, the output of the control system is not the controlled variable, but is an influence on the controlled variable. If the controlled variable were defined as the rotational speed of a motor, the output of the control system would be the torque applied to the motor armature, not the speed. The speed would be classified as an INPUT quantity, because it is this quantity that is sensed by a tachometer, and that can be disturbed by variables in the environment such as friction and loads. We define torque as the output because torque depends only on the output of the control system -- the current going through the motor -- and is not subject to disturbance by the environment. If you are an engineer it will take some effort to reorganize your thinking in this new way, but even in control engineering you might find that there are some considerable benefits in doing so. The normal way of presenting

control processes to students is rather disorganized; the PCT organization brings in a standard approach that often makes control problems easier to solve.

The aim here is to develop some insights into the properties of control systems, not through complexity but through simple examples and hands-on experience. Watching computer simulations work is the next best thing to seeing a real control system work; in some ways it is superior because you have time to see the details of what is going on. Rather than exhort the reader to run these programs on a computer and examine the results, I have decided to make it necessary to do this by not presenting any numerical or graphical tables. You will have to run the program to see what this discussion is about. Perhaps frustration will prove to be an effective motive for actually experiencing this simulation in operation.

Suggestions for standard terminology

A function is a physical device with an output signal the magnitude of which can be computed from the state of its input magnitudes. All functions are true mathematical functions: that is, they may have multiple inputs (arguments) but they produce only one output (value of the function given those arguments). Thus the term function refers both to some physical element of the system and to the equivalent mathematical function that describes the dependence of its output on its input(s) in terms of magnitudes.

A generic control system consists of an input function, a comparator, and an output function. The output of one function generates a variable that is an input to another function. Such information-carrying variables inside the system are called signals. A signal not only represents the value of the function, but serves to carry that value to the input of another function in a different physical location. All signals have a single measure, magnitude. The name of a signal identifies a pathway; the value of the signal indicates the momentary magnitude of the signal carried unidirectionally by that pathway.

The environment model

In the environment of a control system the variables are called quantities. The output of the output function is measured in terms of an effect on a physical variable called the output quantity. The output function in a model of a single control system interacting with an environment is therefore a transducer: its input is a signal while its output is a quantity. The output quantity is always defined so that its magnitude depends only on the output function's value: it is always a single variable. If it has multiple effects in the environment, each of those effects must be separately indicated in a model of the environment.

The input to the control system is another physical variable called the input quantity. The input function senses the state of the input quantity and converts it to a perceptual signal. The input function is also a transducer in a single system-environment model; its input is a physical quantity and its output is a signal. An input function may respond to multiple input quantities.

In the environment, there is a feedback link connecting the output quantity to the input quantity. This link is called the environmental feedback function, or simply the feedback function.


Also in the environment there is a link through which independent environmental variables called disturbing quantities act on the input quantity concurrently with the action of the output quantity on the input quantity. Because the number and kind of disturbing quantities is immaterial, it is customary, when modeling a single control system, to

represent all disturbing quantities and their individual links to the input quantity as a single equivalent disturbance acting through a single equivalent disturbing function.

The control system model

The perceptual signal generated by the input quantity enters a comparator; also entering the comparator is a reference signal, an independent variable. Where possible, the signs of various system constants are chosen so that the reference signal has a positive effect on the comparator while the perceptual signal has a negative effect. The comparator is a function with two arguments and a single value. The output value is represented by an error signal, the magnitude of which is equal to the reference signal's magnitude minus the perceptual signal's magnitude. The error signal enters the output function. Often, as shorthand, we speak of subtracting one signal from another, or adding signals together. What is meant is that the magnitudes are subtracted or added.

The system-environment diagram



DEFINITIONS:

Signals:

sp = perceptual signal
 sr = reference signal
 se = error signal

Functions:

fi = input function
 fc = comparison function or comparator
 fo = output function
 ff = feedback function
 fd = disturbance function

Quantities:

qi = input quantity
 qo = output quantity
 qd = disturbing quantity

THE CONTROL EQUATIONS:

System:

sp = fi(qi)
 se = sr - sp

Interface transducers:

qo = fo(se)
 sp = fi(qi)

Environment:

qi = ff(qo) + fd(qd)

Combined equations:

System: qo = fo(sr - fi(qi))
 Env: qi = ff(qo) + fd(qd)

Setting up a working model

The following discussion assumes that you know a programming language like C, Fortran, Pascal, Modula, or Basic. The actual programming involved is elementary. The student is advised to write the simplest program possible in the most familiar language and experiment with it. The most important knowledge to be gained is a feel for the relationships among variables in a control system, and for the effects of changing various system parameters. There is a temptation to tackle some interesting and complex problem first, but without a strong intuitive foundation for designing more complex systems the most likely result will be confusion and failure. Control systems do many surprising things and the effects of changes in the parameters are seldom what you would initially guess.

The simplest control system to model on a digital computer is one in which all the functions are simple proportionalities except the output function, which is an integrator. Alternatively, the feedback function or the input function can be made into an integrator; however, only one function should be an integrator and the rest should be proportional multipliers. We will use a design with an integrator in the output function; you can experiment with the other possibilities.

In computer programs, integration is summation. Because this is a closed-loop system, integrations do not need to be precise, so advanced methods of numerical integration are not needed. If we make the output function into an integrator, the program step for computing output becomes (in C notation)

```
qo = qo + ko*se*dt;
```


where ko is an integration factor determining how much the output will change on each iteration for a given magnitude of error signal. The constant dt defines the physical time represented by one iteration of the program -- it should be set to 0.1 or 0.01 initially, implying that you should use floating point variables. We will use 0.1 sec.


The central part of a C program for implementing a control system would then be (starting with the computation of the perceptual signal sp):

```
sp = ki*qi;
se = sr - sp;
qo = qo + ko*se*dt;
qi = kf*qo + kd*qd;
```

The disturbing, input, and feedback functions are replaced by constants kd, ki, and kf. For initial experimentation they can all be set to 1. Two variables have to be initialized before the first time this series of steps is used: qi and qo. Initializing to zero is sufficient. Two independent variables must also be set, sr and qd. Using these variables is described below.

The diagram with constants in place of the general functions:





The general flow chart of the program

1. Initialize variables q_0 and q_i .
2. Input or set constants k_i , k_f , k_d , and k_o .
3. Input or set the values of the reference signal s_r and the disturbing quantity q_d .
4. Execute the four program steps above.
5. Plot or print the values of variables of interest.
6. Return to step 4 until the desired number of iterations is finished.
If $dt = 0.1$, 25 iterations will show 2.5 seconds of behavior.

An alternative is to pre-record an array of values for the reference signal or the disturbing quantity or both, and step through this array as the iterations proceed. In that case step six would involve returning to step 3, and step 3 would advance pointers to the arrays of values for s_r or q_d or both. Below we will use still another way of showing the effects of changes in s_r and q_d .

Exploring a control system using the model

```

/* A SAMPLE PROGRAM IN C: */

#include "stdio.h"

void main()
{
float sp = 0.0, sr = 20.0, se = 0.0, qo = 0.0, qi = 0.0, qd = 0.0;
float kd = 1.0, ki = 1.0, ko = 8.0, kf = 1.0, dt = 0.1;
int i;

/* you may put statements here to input values of the k-constants */
printf("\n");
for(i=0; i<25; ++i)
{
if(i > 12) qd = 10.0; else qd = 0.0;
qi = kf*qo + kd*qd;
sp = ki*qi;
se = sr - sp;
qo = qo + ko*se*dt;
printf(
"\x0d\x0a qd=%6.2f qi=%6.2f qo=%6.2f sp=%6.2f sr=%6.2f se=%6.2f",
qd, qi, qo, sp, sr, se);
}
}

```



```
(void) getch(); /* pause to view; press key to exit */  
}
```

In this program, the disturbance remains at zero for the first 12 iterations, and then jumps to 10.0 for the last 13. The resulting table will just fill the screen. You will see the input quantity and perceptual signal rise quickly to 20 units, to equal the reference signal's setting, and then briefly be disturbed when the disturbing quantity changes from 0 to 10. The perceptual signal *sp* will then return within half a second (5 iterations) to the reference value again.

Illuminating experiments with the program

There are several basic rules of thumb that can be demonstrated with this program. The first is the maxim that control systems control their own perceptual signals, not the input quantity and not their own outputs.

First, run the program and note that the perceptual signal *sp*, the input quantity *qi*, and the output quantity *qo* all rise to 20.0 just before the disturbance enters. Note that after the disturbance appears, the output drops from 20.0 to 10.0 units; this is because the disturbance is trying to make the perceptual signal too large; the output automatically drops by the amount needed to bring the perceptual signal *sp* back down to 20.0. The control system clearly does not control to produce a specific output. The output quantity changes as the disturbing quantity changes.

Because the input function is a multiplier of 1, the input quantity *qi* and the perceptual signal *sp* are numerically equal. Change the input constant *ki* from 1.0 to 0.5 and compile and run the program again (to save time you may want to insert some lines to read in *ki*, *kf*, *ko*, and *kd* from the keyboard). Intuitively one might expect that halving the input sensitivity would halve the perceptual signal. But this is a closed-loop system, and that is not what happens.

The perceptual signal still rises to match the reference signal's value of 20.0, although more slowly than before. The input quantity, however, rises nearly to 40. It must do this because we have reduced the effect of the input quantity on the reference signal -- the input quantity must be greater to produce the former amount of perceptual signal. This shows that changing the input function alters the input quantity, but does not alter the perceptual signal's final value. Note also that the output quantity has risen to 40.0 units, as it must do to bring the perceptual signal to the reference value of 20.0. The control system treats a change in the input function as just another disturbance, and alters its output to counteract the change in the perceptual signal. Both the output quantity *qo* and the input quantity *qi* are altered by this change in parameter, but the final value of the perceptual signal *sp* is not altered.

Now restore the input constant *ki* to 1.0, and change the feedback constant *kf* to 0.5. Note that the output quantity now becomes 40 instead of 20 as before, while the perceptual signal still rises to the same value as the reference signal, 20.0 units. This shows that changes in the feedback function will change the final value of the output quantity, but not the final value of the input quantity or perceptual signal.

Finally, if you alter the output integration factor *ko*, you will see a change in the speed with which errors are corrected, but the final states of all the variables are unaffected. CAUTION: keep the product $ko \cdot kf \cdot ki \cdot dt$ less than 1.0. If you make it equal to 1.0 or larger, a computer artifact will be introduced and the system will become unstable. You can try it to see what "unstable" means, but this is not true instability. It's caused by the fact that we're simulating a continuous system on a digital computer.

On an analogue computer this kind of instability would not happen, because "dt" is infinitesimal.

The integration factor is initialized to 8.0. If you use smaller values, the computation will remain stable.

To sum up, when we change the disturbance, the input function, the feedback function, or the output function, the perceptual signal always returns to a match with the reference signal as long as the control system is still working. The only variable that remains under control under all these changes in conditions is the perceptual signal. That is why we say that control systems control their own perceptual signals and not their outputs. They can be said to control their input quantities only if the form of the input function does not change.

You may have noticed that when you alter a function, the resulting change in the final state of the system variables is not at the output of that function but at the input. Changing the input function k_i does not alter the final value of the output of that function, sp , but the input to that function, q_i . Changing the feedback function k_f does not alter the quantity q_i , but the quantity q_o at the input of the feedback function. Causation appears to be working backward. Furthermore, the effect of a change is opposite to the direction of the change. By halving k_i , we caused the final value of q_i to double; by halving the feedback factor k_f , we caused the final value of q_o to double.

These strange reversals of causation are typical of control systems, and account for most of the difficulty people have in understanding how different control systems are from straight-through systems in which causation works in the direction we expect. Of course causation has not really been reversed here, but the feedback effects make it seem that it has.

Relationship of disturbance to output

Restore k_i and k_f to 1.0 and k_o to 8.0, recompile, and run. When the disturbing quantity jumps from 0 to 10.0 in the middle of the run, look at the resulting change in the output quantity. It changes from 20.0 to 10.0. This relationship is not accidental. The effect of the disturbing quantity on the input quantity is just cancelled by the change in the effect of the output quantity on the input quantity. 10 units of disturbance is cancelled by -10 units of change in the output quantity.

Now double the constant representing the disturbance function (change k_d from 1.0 to 2.0). This doubles the effect of a given change in the disturbing quantity on the input quantity. Recompile and run.

As usual, the perceptual signal returns to 20.0 after the disturbance. But look at the output quantity: it drops from 20.0 to 0.0 when the disturbance turns on. Now a 10-unit change in the disturbing quantity results in a negative 20-unit change in the output quantity. The reason? One unit of disturbance now has twice as much effect on the input quantity as one unit of output from the control system. Result: twice as much output is now needed to counteract the same disturbance. That is what happens. It happens not because something "knows" that twice as much output is needed, but simply as the natural result of the operation of the control system.

The "behavioral illusion"

If we consider just the disturbing quantity and the output quantity, we can see that there is an apparent direct relationship between them. It looks as though changing the

disturbing quantity causes the output quantity to change. If we didn't know about the input quantity q_i , we might well think that the system was sensing the disturbing quantity directly, and responding by altering its output quantity. It looks as if a change in the disturbing quantity is a stimulus, which causes a response in the form of a change in the output quantity.

Of course we can see that this relationship holds only because of the input quantity and the fact that the perceptual signal is being maintained at a particular value. The apparent stimulus would alter the input quantity if it were the only influence. But it does not alter the input quantity (for long) because the output changes to have an equal and opposite effect on the input quantity. That is the true explanation of the relationship between the remote disturbing quantity and the output action of the system. The appearance of stimulus and response is an illusion. There can be, of course, true stimulus-response organizations. But an apparent stimulus-response relationship is an illusion when the behaving system is really a control system.

Effect of the reference signal

Right after the bracket following the "for" statement, insert this line of code:

```
if(i>5 && i < 20) sr = 20.0 else sr = 5.0;
```

Recompile and run. Now the reference signal begins at 5.0, and rises to 20.0 halfway through the first part of the run. Then the disturbance is turned on. Halfway through the second part of the run, with the disturbance still present, the reference signal returns to 5.0.

Just follow the behavior of the perceptual signal. You will see that it rises quickly to 5.0, then rises again to 20.0 before the disturbance occurs. When the disturbance rises to 10.0 there is a brief excursion of the perceptual signal which immediately returns to 20.0. Then when the reference signal drops to 5.0 again, so does the perceptual signal.

In fact, the perceptual signal tracks the reference signal in terms of magnitude. The reference signal determines the value to which the perceptual signal will be brought initially, and at which it will be maintained, even if disturbances occur. By varying the reference signal, we can make the control system produce physical effects in the environment that result in corresponding variations in the perceptual signal. Even if disturbances come and go, and even if the feedback function and input function characteristics change (over some range), the system will still produce just the output needed to control the perception at the specified level.

This is why we identify the reference signal with the commonsense notions of intention and purpose.

Dynamic considerations

During the program run, the output quantity of the system becomes whatever it must be to keep sp matching sr , for all combinations of q_d and sr . Immediately after sudden changes in these independent variables there is an error, but the error is soon corrected -- sooner if k_o is larger.

In real environments, physical variables can't jump instantly from one state to another. Normally the changes are smooth; they are also slow enough to allow control processes to begin changing before the environmental changes have gone to completion. To illustrate this, change the line in the program

```
if(i > 12) qd = 10.0; else qd = 0.0;
```

to

```
if(i >= 5 && i < 15) qd = qd + 1;
```

In C this could be shortened but this will work in all languages. The effect is to make the disturbance change smoothly instead of in one jump.

Also, "comment out" the line that alters the reference signal sr, or delete it. Compile and run.

Now the perceptual signal sr rises quickly to the default reference level of 20.0. When the ramp disturbance begins, it rises slightly above 20.0 and remains there until the ramp levels out; then it returns quickly to 20.0 again. You will notice the output quantity changing during the change in the disturbance.

If you reduce ko to slow the system, you will find that the ramp disturbance has a greater effect. The amount of effect that a disturbance has depends on how rapidly it changes in comparison with the control system's speed of error-correction.

By reducing the time interval dt to 0.01 and raising ko to 80, you could make the control system work 10 times faster and reduce the effect of the disturbing ramp by a factor of 10. You would, however, need 10 times as many iterations to cover the same period of 2.5 seconds, and this would not fit on the screen. If you want to see this effect, change the print statement so it prints to the printer, and change the limit of the "for" statement to 250. If you can program graphics, you could plot these variables on the screen and see the behavior in much more detail without using 5 pages of paper.

NOTE: I would appreciate it if programmers versed in BASIC, Pascal, and other languages would write versions of the above program, test them, and transmit them to me for inclusion in an appendix to this primer. Programs should be as generic as possible so as to run on as many computers as possible. I especially need versions for mainframes and workstations. In the future I will simplify the program and make it easier for the user to change parameters from the keyboard. For now, I hope that all programmers and would-be programmers will try the program as it stands and learn from it, even if you have to get some help from a 14-year-old. As this primer evolves I will post revisions and again ask for your help.

Best, Bill P.

Date: Mon Nov 09, 1992 1:12 pm PST
Subject: Marking and models

[From Bill Powers (921109.1100)]

Avery Andrews (921109) --

Your comments on "marking" plus some by Hank Folsom and Dag Forssell have brought my feeling of "something missing" into sharper focus. What's missing in linguistics is a model.

When you talk about marking things in the knowledge base as said or not said, you're IMPLYING a model without saying what it is. Whenever you propose that something gets done in a certain way, you imply a model such that doing this sort of thing is consistent with the properties and constraints of the model. But if you never commit to a particular underlying model, then every time you come up with a new idea you're changing the constraints without knowing it, and perhaps invalidating some aspect of former implied models, also without knowing it. This is the general impression I have had about linguistics -- that the proposals may or may not be consistent with each other in terms of any one underlying model. Nobody can know, because the model is never elucidated.

When I try to explain things like language, I always do it with an eye on the constraints of the HPCT model. I would never propose that perceptions are literally "marked" for communication, because there is no provision in the model for that sort of process. Perceptual signals can't carry anything but the signal; there's no room for a second channel to carry the marking, and there's nothing either to do the marking or to recognize it.

Instead, I would accept that marking indicates some possible process and ask what operations ARE possible within the constraints of the model that would have an effect EQUIVALENT to marking. In this spirit I proposed that nonverbal perceptions are selected as reference signals for systems that produce utterances with those meanings, as per the diagram of a few days ago. The presence of a reference signal and an error is equivalent to marking the nonverbal experience as not said, although there is no explicit marking operation and nothing literally gets marked. The nonverbal experience remains not said -- the error remains -- until some part of the utterance produces a meaning matching the unsaid reference signal, at which point the error signal disappears and the need to say that particular thing disappears.

Now you can play around with all sorts of details in this design, such as how long a perception will persist after the "saying" so that the same thing doesn't have to get said again. There are obviously many details that could be proposed concerning how an error leads to finding a phrase or word that will supply the unsaid meaning, and how the words/phrases get transformed into nonverbal perceptions or meanings. I'm not claiming any particular merits for my sketchy propositions. But if we're to apply HPCT to linguistics, whatever we propose has to operate within the existing constraints of an explicit model. You have to find a way to do what you want done by some use of the principles already in the model.

Only when you are absolutely sure that there is no way to get an equivalent effect within the constraints of the model is there any reason to start changing the model. And then, when you do make a change, you must be sure that it's the sort of change that opens up the possibilities rather than restricting them to a narrower range of phenomena. You don't want an ad-hoc solution that will screw up all the other explanations that the model can provide for other phenomena. You don't want to change the model so that perfectly satisfactory explanations of other phenomena are made clumsier and more complicated. You don't want the model to be more complicated than the phenomena it explains.

The reason I asked about the nature of the knowledge base and how marking might be accomplished is that I was trying to get at your underlying model. What, literally, is the form of the knowledge base in a human being? What aspect of the model perceives this knowledge base, and in what form? How is a marking added, what physical form does a marking take, how is the marking recognized? I don't mean that I want a neural diagram; I'm just asking for a set of rules and definitions that pin the model down to some

particular form, so we can see what its other implications are (other than serving the immediate purpose of explaining how we know something has been said).

If the ONLY function of marking is to explain how we know something has been said, then the concept of marking is empty; it's a dormitive principle. It's just a long way of saying that we know something has been said by means of knowing that it's been said. My proposal, on the other hand, explains this knowing in terms of reference signals and error signals, which have a much broader applicability than just this one phenomenon. This is not an ad hoc explanation; it's a way of finding an explanation within a model that already has properties of general significance and wide application.

Modeling involves thinking in two parallel modes. One mode is concerned with building the model itself -- with laying out its properties, its interacting elements, the fixed rules by which it operates in all situations. The other mode is concerned with phenomena: what do we observe happening, that the model must be able to do out of its own stated organization? The model itself changes far more slowly than our guesses about regularities in observed phenomena. As the model get more specific and more complex, it becomes capable of explaining more phenomena -- but by the same token, it takes longer to see how the model explains particular phenomena.

Sometimes the model really does fail; we can't find a way to restate the phenomena that shows how it would grow out of the model. Then we have to search the model to see why it fails; what is missing, or what is incorrect. But most phenomena seem at first to be outside the model only because we're already interpreting them in terms of a different model. "Marking" is an example. Behind that interpretation are images of the way we would use a pencil-and-paper checklist. But that image is inadequate, because it implies operations and perceptions that are not in the model -- and that are not plausible additions to the model. We can accept that we somehow know what meanings that we want to communicate have and have not been communicated yet, without committing to the checklist model. The checklist model, to be made explicit, would bring in some very complex additions to the HPCT diagram as we try to explain how the creation and use of a checklist are brought about by the model. It would be much better to look for an equivalent process that will accomplish the same end -- not saying again what has already been said, and not failing to say what has not yet been said -- that uses the existing facilities of the HPCT model.

Best, Bill P.

Date: Mon Nov 09, 1992 3:53 pm PST
Subject: Re: language

(penni sibun 921106)

> I agree -- intentions (reference signals) ARE the ONE BIG difference
> between PCT and ALL other models of living systems.
>
>nonsense. intentions are one of the many banes of traditional ai.

Well, maybe you are right. They do talk about intentions a lot.
My impression is that in most cases the intention is modelled as a
command to produce action rather than as a specification for perception.

you're right--their idea of intentions is silly. but it is inaccurate to claim that no other model than pct includes intentions.

>re: below. i can't make much sense of yr cryptic equations. i don't
>think you answered my question. i didn't ask whether you thought
>there were hidden loops in my model. i asked how you could make a pct
>model that accounted for the same phenomena. obviously, i believe
>that the phenomena my model captures are important. i'd like to see
>how you think pct could account for the same process. could you
>please try again?

Ok. This is a very fair question. Perhaps you could describe the
phenomenon that your model "captures".

i guess i'll just repeat the text until you address it. (why are you putting my
techno-jargon in scare quotes?)

> this is a basic sketch of my model: i cast
>the issue as deciding what to say next.
> another way to characterize this might be 1) where are we in the
>linguistic structure (ws) and 2) where are we in the structure of
>``meaning(s)'' (?) we are talking about.
> ignoring the rest of the
>context for the nonce, what is said next is determined both by 1) and
>by 2). this implies that it is *not* the case that the ``meaning'' is
>selected and then the linguistic structure is found to fit it, and the
>next bit of language is produced.

From your discussion with Avery it sounds like
your model (Salix, right?) produces sentences that describe scenes.

oh, dear. you must have forgotten the other language thread. salix (the program that
instantiates some aspects of my model) most emphatically does not produce sentences.
salix produces text in increments that are a word or longer long. the texts can be about
anything that can be encoded in a way salix can understand; the principle criterion is
that the underlying connectivity of the subject matter be explicit or explicable to the
system. salix has produced texts about houses, families, airports, and text style.
examples of salix's output in the first and last domains are at the end of the message.

I think you already have a discreet PCT model -- which
is fine. The references are each of the "descriptions" that should be
produced.

what are you calling ``descriptions''? salix doesn't have any a priori description of
the text it will produce. it has a bunch of stuff it ``knows'' (encoded in the kb), it
has some linguistic tools for expressing that knowledge, and it has a completeness
criterion so it knows when it can stop.

The perceptions are logical variables -- either they are
true or false. The references for each variable is set at true.

no, i don't think that's sufficient.

Then

you go through you inference engine and parser and produce states
of those variables.

no...the inf. engine is incidental--it's not a core part of the machinery. and there's not a parser, under any construal of that term, i think (talk about a word becoming useless once its technical meaning becomes generalized!).

A higher order system is looking for all the reference perceptions to be set to true -- then you are done.

no, a higher order system looks for the completeness criterion to be satisfied, and then salix *can* stop--but doesn't *have* to. can you have that kind of variable in pct?

But I think it would be good to see the data on the phenomenon before judging the value of the model -- control model or not.

i certainly am not interested in your judging the value of my model. if you'd like to look at my data and some of my analysis, please dig out my thesis. what i'd like you to try to do is to assume--for just a few minutes--that i have an interesting and intelligent proposal (viz., what has been said and what there is to say codetermine what can be said now), and work with me to cast that model in pct terms. if we can do that or if we fail in an interesting way, we both might learn something.

[Avery Andrews 921106.1348]

Aside from the one Martin Taylor pointed, out, another difference between language production and ordinary control situations stems from the discreteness inherent in language.

my original mentor in natural language generation called this ``indelibility.'' this has the additional connotation that not only is something said discretely, but it can't be unsaid.

cheers. --penni

=====

nb: each line of text corresponds to one increment of salix's output. the house texts are based on descriptions given by people. the style texts conform to requirements given by a colleague of mine.

house description:

OK
we can start at the side door
and then there's a entrance hallway
then
in the kitchen
there is a window
which is large
and is a picture window
it has two flanking windows
and faces the backyard
and if we're facing that window
on the right is the sliding glass door
and a window
which is small

if we're facing the backyard
on the left is the stove
and a refrigerator
and if we're facing that window
underneath is the sink
then a dishwasher
then there's a living room
it has
a window
which is large
in it
and then there's a short hallway
and
Ann's bedroom on the left
Claire's bedroom on the right
and
a bathroom in the middle
oh yeah
and then there's Penni's bedroom
and a long hallway
that's it

text-style description:

this is a
formal
document
but it still manages to be
clear
non-technical
and
straightforward
and therefore is easy to read

Date: Mon Nov 09, 1992 4:19 pm PST
Subject: Re: language issues

(penni sibun 921109.1300)

i have almost as much trouble constructing enormous msgs. as reading them....

[From: Bruce Nevin (Mon 92112 07:38:16 through Fri 92116 14:05:06)]

>what does meaning have to do with anything? why can't speech actions
>be just like other actions, which presumably aren't mediated by ``meaning''?

I would suggest that all perceptions "have" meaning, that is, any given
perception is linked in associative memory to other perceptions. I'm
not sure precisely what you (penni) mean by the phrase "mediated by
meaning". Let me lay out a bit of what I think is going on.

well, i actually don't believe in ``meaning'' very much. i think it's a systematic error
that people tend to think actions that involve language have meanings and those that

don't don't. i would think that pct would be exactly the kind of approach to expunge ``meanings'' altogether (but then i see y'all using the term).

cheers. --penni

Date: Mon Nov 09, 1992 4:27 pm PST
Subject: means, ends, politics

[From Bill Powers (921109.1430)] Rick Marken (921109.0900) --

Good thoughts on means and ends.

We need to be explicit about one problem with means and ends -- "using the end to justify the means." This becomes a bad thing sometimes -- for example, in adopting the philosophy that it's better to let 20 innocent people go to jail than to let one guilty person go free. The desire for law and order is often used to justify means that involve violation of rights and other cruelties.

The problem here is not so much in trying to reach the goal -- which we can agree is good -- but in forgetting that we have more than one goal at a time. Not only do we want to protect people from criminals, but we want to defend Constitutional guarantees for everyone. The real objection isn't to using the end to justify the means, but in forgetting that the means must support more than one end at a time. This often involves a conflict, and some people don't like the complications involved in finding resolutions in tough situations. So they go with the solution that agrees with their highest-priority goal and rationalize or ignore the other side of the conflict.

This tells you something about a person's priorities. Catching criminals rather than protecting the innocent from harrassment would be the first choice of someone who has power, wealth, and social standing and doesn't anticipate being falsely accused of tawdry crimes. Protecting the innocent from false accusation would be a higher concern for someone who is a member of part of society in which tawdry crimes often occur; someone who has no power or recognized social standing, who can't afford paying court costs and lawyers to fight a false charge, and who looks, dresses, and speaks no differently from the criminals and thus can present no superficial evidence of being an unlikely suspect. There would also be much more concern with apprehending thieves and robbers among those who judge their own worth in terms of their material possessions than there would among those whose main possessions are their friends and their freedom, which are harder to steal.

With regard to your comments on "trickle down" economics:

It has struck me that one of the main discriminators between the clusters of ideologies must be, simply, wealth. The philosophy of trickle-down economics says, basically, let us rich people make all the money we want by any means we can achieve, and some of that will trickle down to the rest of you when we spend it. In the implied economic model, poor people are necessary. It takes fifty or a hundred people at the lower end of the economic scale to make possible the life styles of every one or two who receive most of the income. Given any current level of technology and productivity, there is simply no way that all people can live like rich people. For some people to be rich, many more people must remain poor. The greater the disparity between the bottom and the top, the more people must work for small incomes to allow for a few people to have huge incomes; the more people must becontent with fewer and lower-quality possessions and less fun to allow a few people to acquire many high-quality possessions and play rather than work.

Those with very large incomes naturally want to preserve that situation. They also want to maintain the power needed to protect themselves from the natural efforts of others to get more of the available income for doing less of the necessary work. They promote a work ethic in which self-esteem is coupled to being a good employee who observes the rules and works for the offered wages with gratitude. Rich people can fund lobbies for influencing legislation that would affect their own freedom, their own wealth, and their own power. They are against regulation because meeting regulations would reduce their incomes. They use economic principles instead of morals because the economic principles show how to maximize return on investment, whereas morals might limit doing some of the things that maximize return. The return on investment that they are thinking of is, of course, the return they receive, not that the nation as a whole receives.

So the rich get richer and the poor get poorer, simply because there is no negative feedback to speak of in this social situation. The more wealth a person accumulates, the more power that person obtains to resist restrictions on increasing personal wealth and thus gaining even more power. And the poorer a person gets, the less ability that person has to resist having even more taken away.

The only protections the poor person has to oppose being forced down to the level of subsistence or worse are government and crime. This, I think, is why people at the lower end of the economic scale want more government, and why those at the upper end want less. It is why crime is the most rampant in the poorest neighborhoods. People do what is available for them to do as a way of improving their lives, in their own judgment of what will work.

I am unsympathetic toward those who rail against taxes and government regulation, simply because their own self-interest is so clearly the motive. There can be bad taxes and bad regulations, but that is a matter of quality and appropriateness, not quantity. Those who simply want less taxes and less regulation in general do not have the good of the nation at heart; they are simply trying to preserve what they have regardless of what is good for others. This attitude comes from a narrow system concept in which only some people deserve the best, in which only a few are qualified to tell the rest how to live. It's inevitable that this sort of system concept would create its own mirror image that guarantees endless social conflict. Marx may have had an inkling of this, although he and his successors used a mystical semi-religious theory of history to back up their common-sense predictions. I think that control theory gives us a much simpler way of predicting the outcome of such a "class struggle." When any small group makes life hard for the majority, the majority will eventually get fed up and take whatever steps are necessary to reach their own goals. One doesn't need to be a political pundit or even a control theorist to understand that.

Best, Bill P.

Date: Mon Nov 09, 1992 4:36 pm PST
Subject: Re: Knowledge base; hierarchy & language

(penni sibun 921109.1400)

[From Bill Powers (921107.0800)]

Penni Sibun(921105.1300) -- Avery Andrews (921106.1348) --

I have some questions about the "knowledge base." I gather that what you mean by a knowledge base is a set of statements describing perceptions, rather than the perceptions themselves.

the kb doesn't know from perceptions. it's just part of a computer program! you can think of the kb as made up of propositions--essentially as in 1st order predicate calculus (w/o quantification). a prop. has the form:

(relation x y) where x and y are objects or properties

examples are:

(father-of john penni)
(next-to kitchen livingroom)
(east-of livingroom kitchen)
(west-of kitchen livingroom)
(has-property livingroom-window large)

there's some annotations and bookkeeping hair (like the infamous already-said? bit on objects) but that's pretty much it.

I can see how it would be possible to mark a list of statements as "said" or "not said," but how would you mark the perceptions to which they refer?

there aren't any statements in the kb. the kb contains objects, properties, and relations (for no good reason than that is a classic way of dividing up the a representation of the world that seemed reasonable to me). the kb is structured and interconnected by the way the os, ps, and rs are combined in propositions. while in principle, you, as a human, can ``read'' the kb, a lot of the propositions in salix's kb are never ``stated'', eg, (next-to-doors kitchen livingroom).

In a lot of the proposals offered here by linguists, I get a sense of something missing. PCT itself is about how a brain works (whether we've got it right or not). So I'm always asking "how would the brain do that?"

i think what's happening here is a conflation in our discussion bet. the program salix, which can be described precisely, and my model/theory/whatever about how language works (which includes the brain's role and whatever else). in building a program to demonstrate my theory, i wasn't trying to build an artificial brain, i was building a demonstration that a system with thus-and-so resources can do thus-and-such-a-thing.

I think of a model as showing HOW something is done, even if only in principle;

i think my model *does* show *how* something is done. definitely only in principle, and certainly not at brain-level granularity. my argument that ``what can be said is codetermined by what has been said and what there is to say'' sounds like a model of a process to me! if you don't think so, why not?

Ed Ford failed to say that yearly membership in the CSG for students is \$5. That's one reason the fee for working members is \$45 --

there's a subsidy.

somewhat beside the point, perhaps you can consider subsidising unemployed people too. i'm not suggesting you make an exception for me, but there is a growing class of academics who are poorer than gradstuds cause they have comparable income but are expected to pay full price for everything!

cheers. --penni

Date: Mon Nov 09, 1992 4:41 pm PST
Subject: Re: language

Re (penni sibun 921106)

A small group of us, not all really active CSG'ers have been talking about the question of how to affect a decision maker's perceptions, because the decision maker can only make changes on the basis of perceived truth.

Here is perhaps a beginning of an answer to Penni.

A perception of the real world can be put in terms of a set of propositions believed to be true (Aristotle - or if you want something more modern, Pattern Recognition, Human and Mechanical - Satosi Watanabe, pub John Wiley 1985).

The distance between "your" perception and "mine" is a Hamming distance(*), possibly weighted by value systems for the propositions. Note that you and I can disagree about the magnitude of the weighted Hamming distance because of differing value systems, unless we agree on everything in which case the Hamming distance is zero (perhaps one cause of alienation - whenever I have real row with my boss it's because of different value systems on the same set of propositions).

In this framework, the PCT model has an error signal that is the weighted Hamming distance, and the purpose of communication, diplomacy, or in the last resort war, (von Clausewitz) is to change that distance either by change of perceived truth or change of (even imposing a different in case of war) value system.

Well, there's a whole lot more. We don't have it all sorted out yet, but have promised collectively to do so by mid 1993 in time for a private meeting in Washington DC. But there will be no problem about circulating the talks afterwards, possibly even a book, we don't know yet.

We owe Bill Powers for PCT, C.E. Shannon for Information Theory, Pennti Kanerva for databases having a metric being the the Hamming distance, and Bayes and Goodman for estimates of reliability of perceived truth.

As I say, it isn't finished, it isn't even ready to talk about yet, and perhaps I should have held my peace until we were ready.

Apologies to other members of the group if I've let the cat too far out of the bag, but I think Penni had the string pretty far untied anyhow.

Note by the way, two interesting things. This is not a classical continuous control system - i.e. the PCT concept that you only act on perceptions has deep and far reaching implications beyond the little man, and that there are other formulations of the things

PCT deals with (such as optimisation theory) that become easier to manage when the spaghetti bowl of feedback paths gets too tangled. Conversely, inside every optimisation algorithm there is a PCT system, e.g. Newton's Method.

(*) The Hamming distance is the count of propositions where we disagree. It really truly is a metric (i.e. a distance) that obeys the triangular inequality.

John Gabriel 921109:16:45 CST

Date: Mon Nov 09, 1992 4:56 pm PST
Subject: Re: Language

(penni sibun 921109.1630)

[From Hank Folsen (921109)]

My feeling is that to understand language, rather than starting from the spoken word and working backwards, it will be easier to start from how the human mind functions and work forward using PCT.

if pct were anchored in how the mind works, then the latter w/b an excellent strategy. however, whether pct has a handle on the brain is strictly a matter of faith at this point. one way to strengthen pct's claim to the brain w/b to show it could account for language. but how can we show that? by having a clear picture of what has to be accounted for when we ``account for language'' w/ any theory. in other words, you have to approach a pct theory of language from both directions.

i, on the other hand, would rather approach language from the point of view of interactionism. i'm certainly having to work both sides of the issue myself, in the hopes that i'll ever have a coherent theory.

--penni

Date: Mon Nov 09, 1992 5:28 pm PST
Subject: Re: Marking and models

(penni sibun 921109.1700)

[From Bill Powers (921109.1100)]

Avery Andrews (921109) --

Your comments on "marking" plus some by Hank Folsen and Dag Forssell have brought my feeling of "something missing" into sharper focus. What's missing in linguistics is a model.

When you talk about marking things in the knowledge base as said or not said, you're IMPLYING a model without saying what it is.

i think this discussion of marking is a serious sidetrack. the already-said? mechanism was basically a kludge to allow me to complete my thesis in finite time. i actually think redundancy is fascinating (it happens all the time) but it wasn't my focus. i have a friend in computational linguistics who is doing a thesis on just a class of redundancies (and i may be able to incorporate her work into salix).

in any event, here's all that already-said? ever claimed to do (Locally Organized Text Generation. COINS Technical Report 91-73, Department of Computer and Information Science, University of Massachusetts, 1991. Also Report SSL-91-21/P91-00159, Xerox Palo Alto Research Center, 1991, p. 41) (`{\em}` indicates emphasized text):

Each object [in the KB] has an `{\em already-said?}` field. This field is initialized to `{\em nil}`; when an object has been explicitly mentioned in the text, this field is set to `{\em t}`. The strategies which find something to say next will not be triggered by an object that is `{\em already-said?}`. An object which is `{\em already-said?}` may be mentioned again, but only in relation to another object. For example, in the following text:

- (1) there's Penni
- (2) and then there's Barbara
- (3) who is Penni's sister

In (1), Penni is explicitly mentioned, and marked as `{\em already-said?}`; in (3) she is mentioned in passing, in relation to Barbara.

The reason I asked about the nature of the knowledge base and how marking might be accomplished is that I was trying to get at your underlying model. What, literally, is the form of the knowledge base in a human being?

i'm not even sure that there's a kb-analog in the brain; i bet there is no `{\em marking}`.

can see what its other implications are (other than serving the immediate purpose of explaining how we know something has been said).

`{\em marking}` doesn't have explanatory power. it's a mechanism (and a kludgy one at that, as i said). again, we seem to be getting a conflation of programs and theories.

But most phenomena seem at first to be outside the model only because we're already interpreting them in terms of a different model. "Marking" is an example. Behind that interpretation are images of the way we would use a pencil-and-paper checklist. But that image is inadequate, because it implies operations and perceptions that are not in the model -- and that are not plausible additions to the model.

sigh. right. that's why i'm sorry you got that extra-model interpretation and hope everyone drops it!! the `{\em checklist model}` is not my creation; it's not in **my** thesis!

cheers. --penni

Date: Mon Nov 09, 1992 5:52 pm PST
Subject: Re: Primer for modelers: draft

Hi everybody, I have been listening to your PCT list for a week. I am a Ph.D. candidate majored in Cybernetics and System Theory at George Washington University. I came from China to learn from you. So far my impression of PCT is that it is very close to the idea of Second Order Cybernetics, which has been developed after 1970s by a group of cyberneticians hosted in American Society of Cybernetics. Are you guys familiar with this area? If yes, what are your comments, if not, I will be glad to answer you questions regarding to SOC. The leading scientists in SOC are Heinz von Foerster, Humberto Maturana, Gordon Pask, and Ernst von Glasersfeld. Other active figures include: Stafford Beer, Stuart Umpieby, George Klir and Klaus Krippendorff. Cybernetics itself is concerned with communication and CONTROL in complex systems, and SECOND ORDER means the system being studied includes the observer, (and his PERCEPTIONS,) thus I say that PCT is close to SOC. Could you please offer me some substantial content, not just promotional materials, about PCT, so that I can compare them in detail? I would like to participate your discussion as a cross-disciplinary collaborator rather than a buyer or a user. It is possible that PCT will benefit from other fields like SOC and vice versa.

Thank you and best wishes. Jixuan HU
<jixuanhu@gwuvvm.bitnet OR jixuanhu@gwuvvm.gwu.edu>

Date: Mon Nov 09, 1992 7:26 pm PST
Subject: Study PCT first. Then make revolution.

[From Dag Forssell (921109)]

(penni sibun 921109.1630) in response to Hank Folsom (921109)

>if pct were anchored in how the mind works, then the latter w/b an
>excellent strategy. however, whether pct has a handle on the brain
>is strictly a matter of faith at this point. one way to strengthen
>pct's claim to the brain w/b to show it could account for language.

PCT shows clearly that people attend to that which they are concerned about. To a person primarily concerned about language, who does not know or care about PCT, Penni's suggestion appears to make sense.

Even on this CSGnet, which is dedicated to the study of PCT, there appears to be relatively few people who have studied PCT seriously, as an engineering science. Those who have, recognize and visualize the phenomenon of control. They have studied the book: Behavior: the Behavior of Perception and other works. They have applied the lessons of PCT in their own lives, can see control operate everywhere and have no doubt that PCT has a handle on the brain.

PCT lives up to rigorous levels of testing and verification. Demonstrations with real human beings having to use their brain routinely demonstrate 95-98% correlation between prediction and actual. Does linguistics? Bill noted today the absence of a clear model in the discussions of linguistics. Without a model, you can't begin to test any validity of what you talk about. You cannot possibly live up to any rigorous standards, can you?

I have no doubt whatsoever that PCT can account for language. And well. But it may not account for it on the presently popular terms of a descriptive science of linguistics. To

ask PCT to explain the present state of the art in linguistics is backward. Instead, use PCT to develop a new, rigorous, engineering science of linguistics. To do that, the descriptive linguist must be prepared to think in terms of the physical or engineering sciences, learn PCT first, adopt new and much higher criteria of verification and finally (perhaps hardest of all) be prepared to follow the evidence, even when it may suggest the abandonment of cherished descriptive principles of contemporary linguistics, no matter how famous.

If you learn PCT first, everything will fall in place, your special interest included, whatever it is. If you ask PCT to prove to you what you already think you know (but that may not be so at all), without learning PCT first, you will go round and round but never learn anything from your encounter with PCT advocates, except perhaps that PCT is nonsense and makes excessive claims.

PCT offers the opportunity for a scientific revolution in all the branches of descriptive life sciences. A transition from a maze of descriptions based on each other to an engineering science based ultimately on physics.

PCT is much more than a matter of faith. There is a lot of very hard evidence for those who are willing to look. PCT is knowledge, more certain than most.

Best to all, Dag

Date: Mon Nov 09, 1992 9:23 pm PST
Subject: language & modelling

[Avery Andrews 9211414] (Dag Forssell (921109))

I don't think the lack of good modelling in linguistics by people interested in PCT due to ignorance of PCT--it's due to ignorance of how to build 'essentially connectionist' models (e.g. ones that are consistent with what we know about how neural circuitry works) that actually do the sorts of things that language use obviously involves. I can't do it, Penni can't do it, Bruce can't do it. I suspect that the person who figures it out will be neither a linguist nor a neuroscientist, but somebody with a substantial but maybe not fully professional grade knowledge of both subjects, plus lots of some kind of math. & this person may come along next Thursday or 100 years from now, or never.

At any rate, great sophistication in linguistics is certainly not required to work on this subject - any PCT modeller could start tackling the problem of, say, how to build a device that would hear or read a description of where a place was supposed to be, and then go there. In fact I think that would be an excellent problem to work on for people who think that the linguists on CSGNet are perverse.

Avery.Andrews@anu.edu.au

Date: Mon Nov 09, 1992 11:21 pm PST
Subject: science & engineering

[Avery Andrews 921102.18.13] (Dag Forssell (921109))

>Even on this CSGnet, which is dedicated to the study of PCT, there
>appears to be relatively few people who have studied PCT seriously, as
>an engineering science. Those who have, recognize and visualize the

> ...

>PCT lives up to rigorous levels of testing and verification.

>Demonstrations with real human beings having to use their brain routinely

>demonstrate 95-98% correlation between prediction and actual.

In some areas, but not in others, such as finding the fridge door (since nobody knows how to wire up a fridge-door detector). But the science isn't mature until this and many other problems are sorted out, & it's a big mistake to try to launch an engineering discipline without a clear view of the scope and limits of the extant (as opposed to hoped-for) science (the folks in Eastern Europe are up for quite a few years of picking up the wreckage of such a mistake).

Avery.Andrews@anu.edu.au

Date: Tue Nov 10, 1992 6:11 am PST

Subject: Promoting PCT (posting to biome)

[Bill Silvert 921110] Last week Dag Forssell wrote:

>[From Dag Forssell (921102.1700)]

>

>A disk with updated CSGINTRO.DOC will be postmarked tomorrow for Gary Cziko

>and Bill Silvert. Gary will put it on the net and Bill make it available on

>the server. Pat and Greg have succesfully (I could use it) created an ASCII

>file, uud.scr, which Bill will also put on the server. This allows MCI mail

>(and Bitnet?) netters with DOS machines to get the binary files.

>Instructions in CSGINTRO.DOC. Also, the marken.doc file will be updated,

>complete with ASCII figures. Instructions on how to subscribe to CSG-L have

>been added. Download the CSGINTRO.DOC and pass it around to friends.

The disk arrived and the contents have been put in pub/csg/documents/forssell.

I decided to keep them altogether without editing. Here is the readme file:

[From Dag Forssell (921102)]

Gary and Bill:

Here are four files for the server:

csgintro.doc	Formerly Starter document
uud.scr	
rubberbd.txt	
marken.doc	

Bill: As shown in the csgintro.doc, I propose you place and comment on uud.scr as shown:

programs/source:

uud.scr 54371 ASCII. Compile uud.exe w/ DOS debug. Directions @ end.

In my comment on the net about showing ASCII vs Binary, I anticipated that it would be easy to show ASCII in the comment section as shown here. It is

not a big deal, of course.

You asked why not electronic mail. \$ is the answer. MCI mail postage on the content of this disk (130 kb) would run easily \$10 to each of you. It is priced by the kilobyte for large messages.

I have checked uud.scr with demla.exe. My experiences are reflected in the text.

While there are only a handful of CSGnetters on MCI mail or ATT mail (Ray Jackson) at present, I will do my best to distribute a "demodisk" (available in perhaps a week) to industry. CSGINTRO.DOC will be included. This is one reason for me to spell out how to subscribe and how to access Bill's server. Hopefully, others will find it convenient to pass around copies also.

marken.doc is updated and includes illustrations. Please substitute.
rubberbd.txt is self-explanatory. Minimal comment up to you.
csgintro.doc likewise.

demla.exe currently on the server gives Bill Powers address in Northbrook Illinois. I will mail you an update when I have mastered the archive process so Durango, Colorado is spelled out. I will provide update for dem2a.exe as well. I'll need a week.

Gary:

csgintro.doc should do well for new subscribers and for monthly posting by you.

Best, Dag

--

Bill Silvert at the Bedford Institute of Oceanography
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2
InterNet Address: bill@biome.bio.ns.ca

Date: Tue Nov 10, 1992 7:40 am PST
Subject: Re: Study PCT first. Then make revolution.

[Martin Taylor 921110 10:00] (Dag Forssell 921109)

>PCT shows clearly that people attend to that which they are concerned about.

That's one claim I've never heard on CSG-L before, and I don't believe it to be true. Most of the perceptions that are under good control are not in attention, and I'm not at all sure that "concern" is an attribute of an Elementary Control System. Where emotions and such feelings come into HPCT is a very interesting and hardly touched area.

The only description within PCT that comes close to the concept of "attend" is, so far as I know, based around the core concept that there are far more sensory degrees of freedom than effector degrees of freedom. The discrepancy ensures that there are far more potentially controllable perceptions than are being controlled at any moment. It might be tempting to identify "attention" with the controlled perceptions, but I think it would be wrong to do so. Subjectively, attention seems more associated with changes in which

perceptions are being actively controlled, and thus with the (unspecified) mechanism by which the control shifts among the controllable perceptions.

>PCT is much more than a matter of faith. There is a lot of very hard
>evidence for those who are willing to look. PCT is knowledge, more
>certain than most.

Yes, but there's still an awful leap of faith from there to

>If you learn PCT first, everything will fall in place, your special
>interest included, whatever it is.

Nothing wrong with having faith, but recognize it for what it is.

Martin

Date: Tue Nov 10, 1992 7:47 am PST
Subject: Question on Demo. Promoting PCT.

Thomas Baines (921110.0942 - direct)

Dag,

I appologize for not acknowledging the tape, et al. Thanks. It is being circulated among some interested folk here at ANL. I'll get back to you with comments/requests/etc.

Date: Tue Nov 10, 1992 9:33 am PST
From: Gary A. Cziko
Subject: Intro Document

Dag:

It seems that I never received from your the finished CSGnet intro document. Did you send this to me as you sais you would? Or are you still polishing it up?--Gary

Date: Tue Nov 10, 1992 9:42 am PST
Subject: Language, Interactionism

[From Hank Folson (921110)] (penni sibun 921109.1630)

>however, whether pct has a handle on the brain is
>strictly a matter of faith at this point.

Penni, my understanding and accepting the concepts and claims of PCT for me requires trying to apply them to my world to see if PCT really does work for me, and so is more likely to be true (to me). My interest is in applications of PCT to life in general. (My perception is that PCT is about life in general.) It is working for me so far....

One generally agreed consequence of PCT for believers is that only through force can people make other people do or believe anything. This is because PCT says everyone is an independent control system. I can't sell you on PCT or make you give up your existing belief in interactionism. The only way I can achieve that sort of result is to over time supply your control system with inputs that you can relate to which will gradually build up until your internal comparison engine decides it likes the new inputs better than the old established ones. A decision to change will be made only if the new inputs lead you to have better success in controlling your perceptions than the old ones do. If PCT can't pass that test, the rules of PCT say you will never accept PCT no matter how long or how much information on PCT you absorb. This is because one of the premises of PCT is that control systems continue to do only what works for them in controlling their perceptions. A corollary is that if what the control system is doing stops working or is not working, the system will try at random variations on what has been working, and if that doesn't fix things, it will "reorganize" and try something new, again at random. It is only at this final stage that PCT might get anywhere inside your world.

If you reread Dag Forssell's (921108) post, you may see that Kuhn noticed this aspect of life, as it applies to accepting new theories, but he did not know why it happened. Dag and I are explaining why, in PCT terms that we have learned from others. Just some more background for you.

>i, on the other hand, would rather approach language from the point of
>view of interactionism.

I am ignorant of interactionism and what it is all about. Okay, so I haven't been reading your posts in detail. That is because I am not controlling to learn a lot about linguistics. PCT says that if you are not interested in something (controlling for it, in PCT terms) you will not observe it, or it will go in one ear and out the other (to use my mother's term, as she often does with me).

I am controlling both to learn how to apply PCT and also to minimize the time spent doing it. The following 7 questions are my attempt to apply PCT to a real world situation of interest to us both. The questions are in line with my earlier statements of my ideas about how PCT works, and how new ideas evolve in living control systems.

1. Will you give me a couple of sentences long grossly oversimplified outline of interactionism?
2. Who developed interactionism?
3. About when did they develop it?
4. What did (s)he believe in before (s)he came up with interactionism?
5. What did his/her old belief lack that interactionism offered?
6. How well has interactionism worked for those who use it?
7. Has any independent research backed up interactionism's claims?

If you try to answer these questions about interactionism, I will answer them about PCT.

Hank Folsom

Date: Tue Nov 10, 1992 11:22 am PST

Subject: meaningful or less

[From: Bruce Nevin (Tue 921110 12:55:31)]

(penni sibun 921109.1300) --

OK, Penni, I'll try to keep this short.

>i actually don't believe in ``meaning'' very much. i think it's
>a systematic error that people tend to think actions that involve
>language have meanings and those that don't don't.

I think you're agreeing with what I said, but it seems like you believe you're disagreeing. I said:

(Bruce Nevin (Mon 92112 07:38:16 through Fri 92116 14:05:06)) --

> I would suggest that all perceptions "have" meaning, that is, any given
> perception is linked in associative memory to other perceptions.

Where we differ is this: you would rule meaning out; I urge instead that meaning is pervasive. But we apparently mean different things by "meaning."

You had asked:

>what does meaning have to do with anything?

In the universe of our perceptions (our only knowable universe), meaning has to do with everything. In PCT, meanings are remembered or imagined perceptions associated with a given perception. If we are controlling the given perception, some of these other perceptions arising through associative memory and imagination may also come to be controlled, and in this way affect the actions (behavioral outputs) that result from perceptual control. Among those actions may be included some that we call language.

This is really the crux of my response. If this paragraph is problematic, I'm afraid the remainder will be productive only of red herrings.

You went on to ask:

> why can't speech actions
>be just like other actions, which presumably aren't mediated by
>``meaning''?

I affirm that "speech actions can be just like other actions" but I deny your presumption, now that I understand what you mean by "mediated by meaning" (or I think I do). Traditional concepts of meaning amount only to translating from ordinary language into a metalanguage specialized for the study of semantics (usually with a vocabulary of semantic features and such, and with syntax adapted from mathematical logic). I am not concerned with these concepts and methods, and they have nothing to do with PCT that I know of.

For example, George Boole's "laws of thought" should more aptly have been called "laws for thought." Logic is prescriptive of rigorous thinking, not descriptive of how people do most of their thinking. Logic is almost always used after the fact, either properly, to verify the validity of conclusions reached by other means, or most commonly to rationalize desired conclusions. Meaning is not so well understood as logic because it encompasses the whole universe of one's perceptions. Other approaches to the study of meaning that I know about are limited in analogous ways.

It has been a puzzle to reconcile the self-evident meaningfulness of our experience and actions with the scientific conclusion that the universe is deterministic and meaningless. Understanding that all that we know and can know is the universe of our perceptions (including those perceptions constituting our theories and models) resolves this. PCT appears to be the first well founded basis for a science of meaning.

Bruce bn@bbn.com

Date: Tue Nov 10, 1992 12:24 pm PST
Subject: means-ends;language

[From Rick Marken (921110.0900)] Bill Powers (921109.1430) --

Excellent. Beautiful.

I, of course, did not mean to imply that I thought that the ends justify the means. I was just making a point about control; any controlled result is influenced by the controller and independent disturbances. So, in order to produce a consistent result (end) the controller must vary his/her influences (means) so as to counteract the influences of the disturbance. We all know that controllers can do this because they are organized as negative feedback systems which keep a perceptual representation of the controlled result matching a secularly adjustable reference input (thanks for the GREAT tutorial on PCT modelling, Bill; I'll send a BASIC and pascal version ASAP).

I suppose the article of "faith" (as Martin puts it) of PCT is that ALL controlled results are perceptions that are controlled in this way. This includes complex results that (in Bill's model) are called "programs", "categories", "principles" and "system concepts". It is difficult to deal with these variables quantitatively, so there is plenty of room for argument about what is being controlled (in language, for example) and whether it is controlled (since it is difficult to perform "the test" at these levels).

The fact that higher level variables are controlled has to rely, for now, on somewhat more subjective evidence. For example, we know we control something that might be called "sentence structure" (at least when we write) because we can tell when a sentence seems well constructed and when it doesn't; and we (some of us) will rewrite sentences to make them sound better -- ie. we control something what could be called "sentence structure".

So it is a lot more than faith that leads us to assume that things like programs and system concepts are controlled -- we can experience errors at these levels and we can (sometimes) see people (or ourselves) acting to correct these errors. In order to control these higher level perceptual ends we MUST be able to adjust references for lower level perceptual ends (which are means with respect to the higher level ends). I think this must be considered a fundamental fact of control system operation -- as much as the fact that control systems control their perceptions. This is why I think Ed Ford (921110) is wrong when he says:

>It seems to me that by forsaking our own systems concepts, that which
>we've built over a life time of thinking and experiencing, for the sake
>of what we perceive as social demands or social approval, as More
>states, we lead our country as well as ourselves "by a short route to chaos."

"Forsaking our own system concepts" suggests that we are able to vary our references for our system concept perceptions. For example, instead of having a reference for "being Catholic" we forsake this system concept by setting it to "being liberal Catholic". What

else could "forsaking a system concept" mean in the context of HPCT? Variation in the reference for a system concept (in the current version of HPCT) could happen in 3 ways: 1) as the result of random drift (not a plausible option, really) 2) as the result of reorganization due to intrinsic error and 3) as a result of the disturbance resistance being carried out by higher order systems. The first option suggests that references for system concepts change randomly -- and I can't believe this would work in a control organization. The second two options seem most reasonable -- but in both cases, "forsaking of systems concepts" is DEMANDED in order to maintain control; in the first case, the system concept is forsaken in order to maintain the viability of the organism -- to eliminate intrinsic error; in the second case, it is done to compensate for a disturbance to higher level controlled variables. But in both cases the system concept is a means for controlling another perceptual variable.

So a person must (and will) forsake system concepts in order to control other variables (if system concepts references can be varied; but obviously they can); this kind of flexibility is not a problem; it is a requirement for control. In fact, when references which can be varied are fixed at some level "no matter what" it is a symptom of conflict (as you mentioned later in you post, Ed); and conflict is the enemy of control. When a person controls for beliefs that create conflicts (like the belief that you are a sinner if you don't worship jaweh) then they are limiting their ability to control. I think that people have been far too UNWILLING to forsake some of their system concepts, principles and programs -- much to their own unhappiness (from lack of control due to internal conflict) and that of others (due to lack of control due to efforts to make them have the correct beliefs -- interpersonal conflict).

The control of perception implies flexible means to control perceptions -- unless, as Bill noted, you happen to be rich and well armed. And beliefs -- blind acceptance of certain reference levels for system concepts, principles, programs, etc -- are the enemy of flexibility because they create conflict.

PCT slogan # 2:

Belief is the enemy of control.

penni sibun (921106) --

>i certainly am not interested in your judging the value of my model.

That's fine. I'll stick to PCT.

I will just say that, in the basic sketch of your model, the PCT aspects are quite clear:

- > another way to characterize this might be 1) where are we in the
- >linguistic structure (ws) and 2) where are we in the structure of
- >`meaning(s)'' (?) we are talking about.

From a PCT perspective, linguistic structure and the structure meanings are controlled variables; so they are perceptions with intended (reference) states. That is how your model relates to PCT. I think it would be possible to look in your model and (to the extent that it controls) see what corresponds to the perceptual representation of these variables and their reference specifications. But it seems to me that you are not interested in really trying to see the controlling aspects of your model -- and what it does. So that's fine. You seem to get a little angry at our suggestions (which, at least in Bill's case have been made quite clearly) -- and I don't think there is anything close yet to the kind of data we need to do PCT modelling of language phenomena. I think

language phenomena are VERY interesting but I don't think we'll make much progress on it until someone who really understands PCT starts doing the kind of studies and modelling that needs to be done. I don't have the time and obviously you don't have the inclination. So let's stop fighting. You're happy doing language modeling your way -- go for it!

Best regards Rick

Date: Tue Nov 10, 1992 12:50 pm PST
Subject: re: marking

[From: Bruce Nevin (Tue 921110 12:59:20)] Bill Powers (921109.1100)

Re "marking" something as having been said, I think as Martin has suggested in his Layered Protocol model of communication that we control for a perception of the other party having understood.

The universe of perceptions is not linear, but speech is.

The universe of perceptions is not linear, that is, the ensemble of pathways for shifting attention consecutively from one perception to another that is in some PCT-internal sense "adjacent" forms some sort of network or mesh, not a linear sequence.

Yet when we talk about these perceptions, we must use a linear sequence of words (some of them reduced in form so as to appear as parts of words). This means that we must linearize some particular traversal of the net or mesh of perceptions. Or, commonly, successive traversals with returns to common nodes represented as repeated words (often reduced to pronouns, zero, etc.).

It seems clear that we control for a given traversal of "adjacent" perceptions having been talked about: instead of repeating the traversal in words, we use one of the reduced forms (e.g. "this is why thus and so," where "this" is a reduced form of words already uttered, and corresponds to the same traversal of perceptions that those words referred to.)

Given that we also control for perceiving that a given traversal of has perceptions been understood by the other, such a perception is often grounds for producing reduced forms of words in an utterance, even though their unreduced forms had not actually been previously spoken.

In the analysis of language form (hold your nose Bill, I'm not so far from perceptual control as you might think), you get a more efficient system if you postulate a requirement for word repetition across conjunctions. This provides a basis for reconstructing intervening sentences conjoined under "and" that (ex hypothesi) had been zeroed because they were common knowledge. Thus:

Don't forget your umbrella, it's Wednesday. <==

Don't forget your umbrella, because you use an umbrella when it rains, and the paper said it would rain tomorrow, and you might come back tomorrow if you miss the train, and you sometimes miss the train when you have a late meeting, and you have a late meeting Wednesdays, and

today is Wednesday.

It seems to me that we use the socially learned structure of language as an framework for organizing how our attention traverses our perceptions (the combination of perceptions from the environment with perceptions from memory and imagination). It seems also that the hierarchical organization of perception strongly informs the operator-argument structure of language, as you have suggested, Bill (relationship words for relationship perceptions, etc.). However, that correspondence is muddled by the existence of phrases that correspond apparently to unitary perceptions (frozen expressions, idioms, etc.), diverse perceptions for the same word or phrase (polysemy, homophony), and so forth, but above all because language is a learned, conventional system of social agreements, and the universe of nonverbal perceptions is not.

In our speech, we mark something as "I assume that you already perceive this" by producing only reduced forms of the words for it. We use this in really quite subtle ways to test and confirm (co-affirm) ongoing agreements. This is one not entirely obvious way that the constraints and conventions of language become a very important tool for creating social agreements.

We control for creating and sustaining social agreements. Failure of social agreements is probably occasion for intrinsic error in mammals (possibly in other creatures as well).

Bruce bn@bbn.com

Date: Tue Nov 10, 1992 12:50 pm PST
Subject: PCT

[From Rick Marken (921110.1100) John Gabriel (921110?)

>The distance between "your" perception and "mine" is a Hamming distance(*),

>In this framework, the PCT model has an error signal that is the
>weighted Hamming distance

I think you're talking metaphor, not model. What carries the error signal representing the distance between your perception and mine?

Jixuan

> So far my impression of PCT is that
> it is very close to the idea of Second Order Cybernetics

Keep listening; we will do our best to correct that impression.

Best regards Rick

Date: Tue Nov 10, 1992 1:22 pm PST
Subject: Re: PCT slogans

[Martin Taylor 921110 14:15] (Rick Marken 921109.0900 and 921110.0900)

>While walking with me wife this weekend she challenged me to

>come up with some surprising conclusion derived from PCT that
>might grab the interest of a layperson. Here's what I came up with.
>
>"You can't tell what people are doing by watching what they are doing".

A reference discussed a few months ago:
Vallacher, R.R. and Wegner, D.M. (1987) "What do people think they're doing?
Action identification and human behavior" Psychological Review, 94, 3-15

I guess they aren't laypersons, but they aren't PCT persons either. They would agree with
PCT slogan number one.

>PCT slogan # 2: Belief is the enemy of control.

I would have thought belief to be the essence of control. Belief, to me, is a short word
for "perception of the current state of whatever is being controlled for by this ECS."
Alternatively, belief is the concatenation of possibly many percepts of the same level.
I don't understand the sense in which Rick means PCT slogan #2. It is obviously
different from what I mean by it.

Martin

Date: Tue Nov 10, 1992 2:06 pm PST
Subject: Hamming it up

(John Gabriel 921109:16:45 CST) --

> A perception of the real world can be put in terms of a set of
> propositions believed to be true (Aristotle - or if you want
> something more modern, Pattern Recognition, Human and Mechanical -
> Satosi Watanabe, pub John Wiley 1985).

Henry Hiz wrote some interesting papers on this so-called "aletheic theory of truth" in
the 1960s or '70s. The direction I took it in my MA thesis at Penn (1970) was to build
an acceptability model of propositions rather than a truth model. Those
operator-argument combinations actually encountered in utterances (some of them
sentences, some not, Penni) accepted by my desired peers are ipso facto acceptable.
Others are acceptable by analogy over the classifier vocabulary of the subject matter
domain. The simplest case: (a) substituting a classifier word Ncl for a word actually
occurring in a given utterance and (b) substituting for the classifier word Ncl each of
the words that is acceptable in an utterance of the form N is an Ncl. This corresponds
to some of the sublanguage analysis work of Sager, Grishman, Hirshman, and others. The
more remote the analogical extension, the lower the acceptability or the expectation of
encountering or producing a particular operator-argument dependency. From these
differences follow the enabling conditions for the reductions that make up most of the
messiness of language.

> The distance between "your" perception and "mine" is a Hamming
> distance(*), possibly weighted by value systems for the propositions.
> Note that you and I can disagree about the magnitude of the weighted
> Hamming distance because of differing value systems, unless we
> agree on everything in which case the Hamming distance is zero

The only way that you can arrive at a Hamming distance of zero in this sense is if you (the observer arriving at the Hamming distance of zero) occupy an external point of view able to look inside each of the two parties and inspect their perceptions from the outside. In general, the Hamming distance can be no more than the count of propositions that one person affirms and with which this person perceives that the other disagrees. There are ways of negotiating and testing agreements, of course, but they are never complete.

Also, in PCT you are dealing with perceptions, not with propositions. (Propositions are also perceptions, that is, the linguistic perceptions whose control results in speech, but it is nonverbal perceptions that I assume are the object of agreement and disagreement here.) The division into propositions and "weighting by value systems" is probably at least in part an artifact of this confusion. If you are representing perceptions in your model by propositions, then it isn't a model of perceptual control (a PCT model).

> In this framework, the PCT model has an error signal that is the
> weighted Hamming distance,

You mean in your model there is an elementary control system with a perceptual input signal corresponding to the number of propositions about which there is disagreement, weighted in some way according to the values associated with the propositions in a value system, and a reference input signal of zero, such that any perceived disagreements result in a non-zero error output signal, right?

>and the purpose of communication,
> diplomacy, or in the last resort war, (von Clausewitz) is to
> change that distance either by change of perceived truth or
> change of (even imposing a different in case of war) value system.

This says that people control for agreement. Or am I missing something?

Bruce Nevin bn@bbn.com

Date: Tue Nov 10, 1992 2:19 pm PST
Subject: Re: Fundamentals of HPCT, Version 0

[From Oded Maler 921110]

* [From Bill Powers (920926.2200) "Fundamentals of HPCT, Version 0"]
*
* D-12. Conflict, if important in the sense of affecting the critical
* variables in the systems that have lost control, will lead to
* reorganizations that escalate the conflict, bringing both loop gains
* into the high region and increasing the outputs to their maximum. This
* will continue until one or both organisms dies, or until one organism
* reorganizes in a way that eliminates the conflict altogether.
*
* D-13. Therefore the normal steady-state interaction between organisms

* will be one that involves no conflict.
 *
 * D-14. Interactions without conflict require that independent organisms
 * in close contact control perceptual variables that are linearly
 * independent of each other. This is the solution to conflict that will
 * occur, through reorganization, if both organisms survive. It is the
 * only solution, other than breaking contact, that can persist over
 * time.
 *
 * D-15. When there are no conflicts, each organism controls variables
 * that are independent of the other's controlled variables. The actions
 * of each organism can amount to disturbances of variables or parametric
 * disturbance of the other's control systems. These disturbances,
 * however, remain within the range where each system can continue to
 * control all its perceptions by varying its own output actions. Each
 * organism continues to maintain its own critical variables near their
 * reference levels, independently of the other.
 *
 * D-16. The shared environment will then come to a state in which it
 * simultaneously satisfies all the reference conditions in both
 * organisms. This can involve a close intertwining of physical
 * relationships, because for the behaving systems, "environment" is all
 * that is not CNS or reorganizing system. It includes, therefore, the
 * bodies of both systems, not just the world between the systems.
 * -----

I wonder how much of the above holds if you replace "organism"
 by ECS.

*D-14. Interactions without conflict require that independent organisms
 * in close contact control perceptual variables that are linearly
 * independent of each other.

Can you elaborate a little on the coexistence of "close contact" in one hand "control of linearly independent variables" on the other. I can see that if 2 systems perceive the same point in the plane, one controlling for the x dimension and one for the y, then they can achieve their goals simultaneously, but in what sense are they then in "close contact"? Since they are controlling for lin. indep. perceptual variable, they are virtually disconnected from each others (from their perceptual point of view). Only some hypothetical 2-dimensional CEV makes them connected.

--Oded

Date: Tue Nov 10, 1992 3:33 pm PST
 Subject: Linguistics, misc

[From Bill Powers (921110.1200)]

Penni Sibun (921109.1300) --

>well, i actually don't believe in ``meaning'' very much. i think
 >it's a systematic error that people tend to think actions that
 >involve language have meanings and those that don't don't.

PCT and I agree with at least your second sentence. It isn't actions that have meanings, whether they're verbal or other motor outputs. As Bruce Nevin pointed out, in PCT (and elaborations thereon) meanings are perceptions evoked by other perceptions. If you perceive the word "green", its meaning is the color that comes to mind. If you perceive the color green on a traffic light, its meaning is a perception of your car starting to move (in that context). This translation between one perception and another associated one is our proposition for how the folk concept of meaning works.

As to doing without meaning entirely, I think that would pretty much put an end to communication.

>a higher order system looks for the completeness criterion to be
>satisfied, and then salix *can* stop--but doesn't *have* to. can
>you have that kind of variable in pct?

Sure. It's what we call a "single ended" or "one-way" control system -- one that considers a departure of a perception from the reference signal in only one direction to be an error. This is like controlling for how close to the edge of a cliff you like to perceive yourself. When you're closer than the reference distance, there's a big error and you pull back. But if you step back so you're farther from the edge than that amount, the control system doesn't act to bring you closer again. One-way control.

Actually all simple neural control systems have to be one-way systems, as I model them in my head. The reason is that comparators are made of neurons with one inhibitory and one excitatory input, the excitatory one being the reference signal. You get an error signal only when the perceptual signal is less than the reference signal, leaving some effects of the reference signal uninhibited. You increase your action until the perceptual signal completely inhibits the effect of the reference signal, leading to zero error signal (the output of the neuron). When the perceptual signal is more than large enough to completely inhibit the reference signal, you still get zero error. Neural frequencies can't go negative -- less than zero impulses per second.

To get a two-way neural control system you need two comparators, one of which has inhibition and excitation swapped relative to the other. This implies that control actions that pass through a zero or a neutral state will not, except by accident, be completely symmetrical. If you measure the parameters of control actions, separating the actions for positive errors from those for negative errors, you should find a slight difference in the parameters. We haven't tried that yet.

>... what i'd like you to try to do is to assume--for just
>a few minutes--that i have an interesting and intelligent proposal
>(viz., what has been said and what there is to say codetermine what
>can be said now), and work with me to cast that model in pct terms.

For this you must propose a model (my meaning). The model would be a proposal as to HOW what has been said and what is to be said do this codetermining. You may actually have built some of this model into your program, without identifying the process implied by the program manipulations in PCT terms.

Suppose that in the knowledge base there are elements called John and Penni. Whether you mean the person or the name of the person is immaterial. In PCT terms we would say that the program perceives things, names them, and stores the names in memory structures. The perceptual signal would be whatever is used in the program to stand for the memory structure -- a 32-bit address, for example, although in a closed-loop system any

arbitrary but unique symbol would do (it won't be necessary in a closed-loop system to work backward from the symbol to the thing it indicates).

Also in your knowledge base you want some relationship perceptions to exist -- for example (father-of John Penni). This implies a general perceptual function of the form (father-of x y) where x and y are any two inputs from the lower level elements. I don't know how you'd design such a perceptual function; presumably it would rely on multiple sources of evidence other than just the names John and Penni. There would have to be something about John and something about Penni and something about their relationships that would reveal that the first element is the father of the second element -- or is not.

Of course you could let this process occur behind the scenes or under the table, and just deposit the specific statement (father-of John Penni) in the knowledge base. In that case you would be relying on processes in the programmer that identify the father-of relationship, but not specifically representing those processes in the model.

The same could be done at the level of the individual words. Instead of proposing some process by which a set of sensory inputs is identified as "John" or "Penni," you could just put those terms in the knowledge base, letting the processes of perceiving and naming take place outside the model.

This takes us only the first step toward a PCT model of how what has been said and what is to be said codetermine what is to be said next. We have been trying to determine how something can be known to be said. We have yet to determine how the system can know what remains to be said, so more parts of the model can be constructed. Then we will have to make some proposals as to what "codetermination" is to mean. And finally, we will have to close the loop, so that what is finally said can be made to satisfy all the conditions we have set up in the earlier parts of the model.

I would be very surprised if the answers to many of these questions were not already embodied in the steps of your program. To unravel the PCT equivalent of the model that you are already using (if such exists), it would be necessary to analyze the program itself to see how one kind of information is extracted from other information, to see how comparisons are made, and to see how the results of comparisons lead to actions. I won't dispute that the outcome of all these processes will be a codetermination of what is to be said next by what has been said and what remains to be said. But we won't have a model to account for that phenomenon until we have taken your program apart and seen just how this codetermination is brought about.

The best person to do that, I suggest, is you. If you don't want to take the time off to learn the fundamentals of PCT, then it's not likely to get done. It's up to you to judge whether the result would be worth the effort.

Dag Forssell and Hank Folson --

PCT has been shown to make very accurate predictions for some very simple tasks. I get extremely uncomfortable when I see public claims that the PCT model is 99.999% accurate. That is not just misleading, it's a lie. Saying things like that will make people think that PCT can predict what people will do with that kind of accuracy in all situations. This is simply not true. There are years and years of experimentation and model-building to do before we can know how well PCT will apply in areas such as language, social interactions, learning, and so forth -- all the fields where PCT has not yet been systematically applied or tested.

If you want to sell PCT that's fine. If it has engaged your enthusiasm, wonderful. But please tone down the hype. That doesn't further the cause. Sorry about the bludgeon, but I want to get the donkey's attention.

Jixhuan Hu (921110) --

Welcome to CSGnet! I think that before you start making comparisons between our uses of control theory and cybernetics, you should read some of the introductory material. It takes most people more than a week to realize what we're talking about. You won't find much in the kind of cybernetics you were mentioning that is like what we do.

Best to all, Bill P.

Date: Tue Nov 10, 1992 3:41 pm PST
Subject: Re: PCT slogans

Martin Taylor (921110 14:15)

>A reference discussed a few months ago:
>Vallacher, R.R. and Wegner, D.M. (1987) "What do people think they're doing?
>Action identification and human behavior" Psychological Review, 94, 3-15
>
>I guess they aren't laypersons, but they aren't PCT persons either. They
>would agree with PCT slogan number one.

But would they understand it?

Could you give a quick summary of the point of their article (which I remember reading and finding irrelevant -- though I always have hope in my soul).

>I would have thought belief to be the essence of control. Belief, to me, is
>a short word for "perception of the current state of whatever is being
>controlled for by this ECS."

Well, I have a lot more confidence in my perception than that; what you described is what I would call "knowledge". I think of a belief as a potential perception -- like "the moon is made of green cheese" -- which may or may not be subject to test. In the control model, I would identify beliefs with reference signals -- not perceptions. Beliefs are typically experienced as imaginations -- the reference signal played right back up as the perceptual signal. We probably develop these "belief references" because the imaginations they produce satisfy other, higher level references; the belief in god, for example, probably makes the believer feel good because it contributes to the perception of their "self" as something special. So beliefs can be helpful -- but they can also lead to conflict (since beliefs often require, as a matter of principle, that they not be revised). This kind of belief (as many pain in the ass athiests like myself are wont to point out) is placing a real bug in the hierarchy. Of course, it is no bug if you can manage to stay away from situations which "push" on the conflict; situations where the person with the bug feels "out of control".

Best regards Rick

Date: Tue Nov 10, 1992 4:19 pm PST
Subject: Hamming distance; belief;association;fundamentals

[From Bill Powers (921110.1430)] John Gabriel (921110) --

Let me echo Rick Marken's and Bruce Nevin's comments on Hamming distance.

I think your concept will be most useful in a model of the way an individual compares the meaning of a communication that is under construction with the intended meaning. This has come up repeatedly; Avery Andrews is particular has commented several times that implementing my idea of closed-loop sentence construction is difficult because of not knowing how to judge an error in meaning, and how to transform the error into an appropriate kind of change in the sentence or utterance (Penni has made me self-conscious about saying "sentence").

But I agree that it's not very useful as a judgment of the difference between two people's perceptions, either of sentences or of meanings. A comparison has to be made inside one of the systems. Two people may perceive different differences between their perceptions (ugh, what a sentence). We can't use that old disembodied omniscient observer ploy. The FACT that two people perceive the same thing is unknowable unless someone knows it.

Martin Taylor (921110.1415) --

Rick: >Belief is the enemy of control.

Martin: >I would have thought belief to be the essence of control.
>Belief, to me, is a short word for "perception of the current state
>of whatever is being controlled for by this ECS."

You're using belief in the sense of perception; Rick was using it in the sense of a reference signal. I think that Rick's use has to be qualified to read:

Unchangeable belief is the enemy of control.

A fixed reference signal anywhere in the hierarchy is bad news for the organism. If it's not at the highest level, it makes the associated control system unusable by higher systems -- they can't vary the reference signal when they need to.

I think system concepts -- or whatever the highest level is when posterity reads this -- have to be a special case. The reference levels at the top level can't be varied freely by a higher system because there's no higher system. It's my impression that system concepts are extremely hard to change. Hard to learn, and hard to change. They come along, they grab you, and then they have you. I don't know how that works. But you can see that it DOES work just by looking at Islam, or any other system concept shared by a lot of people (or so they believe). I don't think anyone can say "I'm bored, I think I'll change a system concept." I don't think anyone changes a system concept without turmoil, pain, and confusion, and the result is not likely to be predictable. This is why PCT is slow to be appreciated as a new system concept. Much easier to see it as a couple of new principles, or a set of clever calculations.

Bruce: > In PCT, meanings are remembered or imagined
>perceptions associated with a given perception.

Martin: >Sorry? What construct in PCT is an association?

There isn't one. We're fooling around with the phenomenon of association, which I think we have to admit exists, trying to find a good place in the model for it. Your suggestion that it could be part of perceptual functions is good. It could also form a link BETWEEN perceptual functions, which is pretty much what Bruce is thinking of, I think. I agree that there may be other more PCTish ideas that can explain SOME phenomena of association. But I don't see anything yet that can explain my being reminded of a Volkswagen by a certain engine sound. Nothing, that is, but content addressing of memory. I think of association as a phenomenon waiting to be put in the model.

There could even be a specific kind of perceptual transformation that involves inputs of one kind and a perception of another kind, where the link is memory association and not an ordinary computing function. Keep thinking of ideas. Maybe the right one will show up.

Oded Maler (921110) --

RE: Fundamentals of PCT

>I wonder how much of the above holds if you replace "organism" by ECS.

Some but not all, because I think that section was concerned with social principles. In social situations, there is no superordinate control system that senses and sets reference signals for lower level control systems. Inside one organism, a higher-level system is responsible for conflict; it is specifying incompatible reference signals for ECS inside one person. As a result, the conflict can be resolved by reorganizing the higher-level system. This isn't possible in a social conflict, where there is no higher-order system in charge of both persons.

>Can you elaborate a little on the coexistence of "close contact" in
>one hand "control of linearly independent variables" on the other.
>I can see that if 2 systems perceive the same point in the plane,
>one controlling for the x dimension and one for the y, then they
>can achieve their goals simultaneously, but in what sense are they
>then in "close contact"?

When the controlled variables are really orthogonal, close contact doesn't exist, of course, even though the two systems may be controlling something associated with the same object. But that's an end-state; normally we don't start out interacting with others orthogonally, and even after learning we seldom achieve perfect orthogonality. I meant the term "close contact" to be a contrast with the situation where completely different objects are involved, so no attention has to be paid to orthogonality.

Put it this way: "close contact" implies the likelihood of conflict, which can be resolved by achieving orthogonality of controlled variables.

It's always nice to see evidence that someone understands what I'm talking about even when I don't say it very well.

Bruce Nevin (92110.1259) --

>It seems clear that we control for a given traversal of "adjacent"
>perceptions having been talked about: instead of repeating the
>traversal in words, we use one of the reduced forms

Sounds good. Do you agree that we also control for unwanted meanings? I.e., some traversals lead to words that have unwanted meanings that stick out enough to cause us to

switch to a different traversal, or to insert a specific denial of the unwanted meaning, or to loop back and repeat part of the traversal in different terms.

>It seems to me that we use the socially learned structure of
>language as an framework for organizing how our attention traverses
>our perceptions (the combination of perceptions from the
>environment with perceptions from memory and imagination).

Yes, I agree that when we think in words (not always the case) we use the socially (well, individually) learned structure as a way of guiding the traversal of meanings -- for creating a scenario to go with the words. This also applies to what happens when we listen to another person speaking, or read, doesn't it? The different in listening or reading is that we have to supply in imagination a lot of the details that the words don't evoke. We don't often misunderstand ourselves, but it's easy to do when listening to others.

>Failure of social agreements is probably occasion for intrinsic >error
in mammals (possibly in other creatures as well).

I can think of lots of ways in which such failure would lead to intrinsic error without the failure itself being an intrinsic error.

Best to all, Bill P.

Date: Tue Nov 10, 1992 5:24 pm PST
Subject: linguistics and modelling

[Avery Andrews 921111.1015]

Pre-molecular genetics might be a good analog of what I suspect will be the relationship between linguistics and `core' PCT for some time to come. People knew quite a lot of useful things about genes, without knowing exactly what they were. This knowledge was genuinely scientific, and essential to identifying them correctly, but of course things were much better when people actually knew what genetic material was.

Avery.Andrews@anu.edu.au

Date: Tue Nov 10, 1992 5:45 pm PST
Subject: Re:dropping babies second try

[From Dick Robertson 921110 repeat of 921105]

Thanks to Gary and Bill for clarifying what happened with my last post. I often try to use this automatic reply program when I'm in a hurry, because me teaching load is so heavy this term that I hate to take the time to write and then download. But I'm out of here next May (Yippee) and hope to have more time to actually keep up with all these great new developments.

Anyhow - What I was saying is that I got the baby-dropping example (actually it was for a test on chapter 5 in the text book) from a real event. Many years ago someone was telling a group of us friends how a relative of theirs got her face all skinned. She had been carrying her baby over some ice and slipped, and fell right on her face while holding the baby up out of the way. I had speculated that it was not a violation of reaction times being faster at lower levels but rather, that she already knew there was ice when she

started out and hence some higher level (principle probably?) had set some command to the ongoing program like, "Don't drop the kid no matter what." Thus, I assumed, the lower orders that would be involved in flinging out one's arms when losing one's balance were "essentially locked" by the commands filtering down from the program. I didn't think that part of my speculations through very deeply to see if that was the best explanation, since I have seen other episodes of similar happenings, I believed that I had observed other instances of lower order controls being "locked," like that. I hope Bill will comment on whether there is a fuller explanation.

--New topic -- When I leave next spring our department will be looking for a replacement (maybe 3 because we haven't filled several past vacancies) and if anybody know of a new HPCT-wise PH D in perception/learning, clinical or neuro- psych who is looking for a teaching post in a mainly undergraduate program-- please send'm this way. It would be great to think that the two rival intro- general psych courses here (one traditional versus one fully PCT) would be con- tinued.

Best Dick Robertson

Date: Tue Nov 10, 1992 5:46 pm PST
Subject: Re: meaningful or less

(penni sibun 921110.1600) --

[From: Bruce Nevin (Tue 921110 12:55:31)] (penni sibun 921109.1300) --

just to add one more cycle to this, i think we're disagreeing even less that you think.

Where we differ is this: you would rule meaning out; I urge instead that meaning is pervasive. But we apparently mean different things by "meaning."

the concept of ``meaning,' ' or anything else, is useful in a theory only insofar as it has explanatory power. you seem to be saying here that meaning is everything--is inherent. to say that something is everything isn't much different from saying it is nothing: what good does it do to say ``that has property x' ' when everything has property x?

You went on to ask:

> why can't speech actions
> be just like other actions, which presumably aren't mediated by
> ``meaning' '?

I affirm that "speech actions can be just like other actions" but I deny your presumption, now that I understand what you mean by "mediated by meaning" (or I think I do).

the rest of your discussion was interesting, but i didn't see it as addressing your point. it's not clear to me whether you think speech actions are distinguished actions or not. (my argument is that they are not distinguished (in any nonincidental way)).

cheers. --penni

Date: Tue Nov 10, 1992 8:28 pm PST

Subject: Closed Loop

[From Dag Forssell (921110)]

Just a note to let Greg and Ed know how much I enjoy the Closed Loop and Ed's card.

I am noticing how skillfully Greg chooses a (later, as I recall) piece by Ed to introduce the thread. The whole piece flows well and holds interest. Thanks, Greg.

Ed, your card is a treasure of specifics. Soon you will have to publish it on a plastic card with built-in magnifying glass. The section on setting standards reflects a heavy emphasis on counseling juvenile delinquents? You certainly operate against a background of threats of violence. Your accomplishment is to reduce the level of coercion substantially and teach the juveniles to function better, which in turn reduces the level of coercion even more.

Best, Dag

Date: Tue Nov 10, 1992 8:29 pm PST
Subject: there's more to meaning than perception

Although meaning is obviously grounded in perception, I think people may be underestimating the subtlety of the connection. Consider the meaning of the NP 'the Iliad'. There are many different kinds of perceptions that this can be hooked up to. In the benign environment of your local library, you can ascertain that a book is the Iliad by noting that it has 'Iliad' written on its spine, or by opening up and seeing that the text starts out 'Sing, Muse, of the wrath of Achilles', So a wide range of different perceptual functions will suffice distinguish copies of the Iliad from other things. On the other hand, you can imagine more treacherous worlds, where books say on their covers that they are the Iliad, but have something completely different inside, or have all sorts of omissions and interpolations in the text, so that some expertise (e.g. a large and sophisticated perceptual function) is required to distinguish good copies of the Iliad from bad ones (the Ancient world was like this, before the Alexandrians cleaned up the text).

But regardless of what kind of perceptual function you may have to recognize copies of the Iliad, I am in the wrong if I sell you something on the basis that it is a copy of the Iliad, & it isn't acceptable as such to an expert (if I do this on purpose, I'm a crook, otherwise a jerk). So meaning depends on perceptions, but not necessarily perceptions of any of the people actually involved in a linguistic communication where the meaning figures.

Avery.Andrews@anu.edu.au

Date: Tue Nov 10, 1992 10:14 pm PST
Subject: Re: Hamming it up

Thankyou for your useful comments. Here are my best attempts at answers.

Last first - I had an assumed context, if you like, the relation between a decision maker setting policy, and their staff. Problem - how to best support the decision maker. The staff want to get the best information to the decision maker as soon as possible. If the staff are good collectors, then this comes down to making sure staff and decision maker

are singing from the same sheet of music (same Aristotelian Propositions) and optimising channel bandwidth - there are some things rather like maximum power transfer and impedance matching, but a bit beyond what seems to be in the literature. In these circumstances staff will need to control for best understanding by decision maker, and to flag any new Aristotelian Propositions not yet in the decision maker's reference set, and develop Hamming distance like associations with things already known to decision maker. Essentially the game of 20 questions (20 bits can distinguish between elements of a quite large ensemble).

It's slipped my mind what the other two questions were, and being logged in on a dumb terminal I can't open a window to reread yr. mail. Can I finish my reply when I come back from out of town. Early next week I hope.

I didn't intend to suggest that people ALWAYS control for mutual understanding. At times deception is wanted - see for example WWII where it was worth many British and American lives to deceive about the place of landings on D-day.

But if speech is going on in order to communicate, then I think people DO control for understanding, and that was the assumed context of my reply to Penni. And Penni's statements from the K-B seem to me first cousins of Aristotelian propositions. And the distance between beliefs and experience does seem to me to be a Hamming distance, probably with weights - if you have just arrived in a strange town, you are probably more concerned about finding your hotel than the dairy.

Sorry - there's the 'phone. Have to go. Will continue later.

John 921110 23:01 CST

Date: Tue Nov 10, 1992 10:35 pm PST
Subject: Re: Hamming it up

[John Gabriel (921109:16:45 CST)] --

Can't resist - a new insight from Bruce's mail. If Penni's or my KB really is a space of 2^N truth values for N Aristotelian propositions defining a mutually agreed universe of discourse, and this space is sparsely populated by instances (not hard for say $N=256$), then one can triangulate a new instance by its' Hamming distances from old ones. BIG question - how do we put the benchmarks in?? I think by a set of shared experiences - sort of like the light signals of relativity that define coincidences in space time - these are a physicist's abstraction for the shared experiences that we all have as a society, and which are the basis for communication. Another question - how much is the lack of acceptance for PCT by many cognitive scientists caused by lack of experiential benchmarks in control theory, how much by badly triangulated ones, (I remember a quite well known ecologist claiming quite seriously that animals eat entropy - terrible confusion between Boltzmann entropy and Shannon entropy), and how much simply by different value systems. This is an extension of a comment by Bill Cunningham about my saying that most of my fights with my boss were about different values. He said No! many of the bigger fights are about Aristotelian propositions, i.e. missing or badly triangulated benchmarks.

John Gabriel 921111 00:02

Date: Wed Nov 11, 1992 6:50 am PST

Subject: *INTRO TO CSGnet*

[from Gary Cziko 921111.0515 GMT]

Here is the introduction to CSGnet prepared by Dag Forssell. I will post this monthly to CSGnet and send it directly to all new CSGnet subscribers when they subscribe.

I will use the same subject description each time (*INTRO TO CSGnet*) so you may just ignore it if you already know what CSGnet and PCT are all about.

If you have suggestions for changes and/or additions, please send these directly to Dag at (e-mail address: 4742580@mcimail.com).

Thank you, Dag, for pulling all this together for us and for the subscribers-to-be.--Gary Cziko

INTRODUCTION TO THE CONTROL SYSTEMS GROUP NETWORK (CSGnet)
AND TO THE CONTROL SYSTEMS GROUP
Prepared by Dag Forssell 921102.

This introduction provides information about:

Why you might want to join the CSG and read CSGnet
Our subject matter: The control paradigm
The evolution of the control paradigm
Demonstrating the Phenomenon of Perceptual Control
The purpose of CSGnet
CSGnet participants
How to ask effective questions
The Control Systems Group
Subscribing to CSGnet
How to obtain text and program files
Literature references

WHY YOU MIGHT READ THE CSGnet

If you have an interest in how people and organisms work. If you are dissatisfied with the explanations in many of the "soft" life sciences and would like a more rigorous approach.... If you insist on thinking things through for yourself....

OUR SUBJECT MATTER: THE CONTROL PARADIGM

Human control is the primary subject of CSGnet, but all forms of control are game. Here is a brief introduction by the primary creator and promoter of the application of the control paradigm to living systems, William T. Powers:

There have been two paradigms in the behavioral sciences since the 1600s. One was the idea that events impinging on organisms make them behave as they do. The other, which was developed in the 1930s, is PERCEPTUAL CONTROL THEORY (PCT). Perceptual Control Theory explains how organisms control what happens to them. This means all organisms from the amoeba to humankind. It explains why one organism can't control another without physical violence. It explains why people deprived of any major part of their ability to

control soon become dysfunctional, lose interest in life, pine away and die. It explains what a goal is, how goals relate to action, how action affects perceptions and how perceptions define the reality in which we live and move and have our being. Perceptual Control Theory is the first scientific theory that can handle all these phenomena within a single, testable concept of how living systems work.

William T. Powers, November 3, 1991

THE EVOLUTION OF THE CONTROL PARADIGM

The PCT paradigm originates in 1927, when an engineer named Harold Black invented the negative feedback amplifier, which is a control device. This invention led to the development of purposeful machines. Purposeful machines have built-in intent to achieve consistent ends by variable means under changing conditions.

The discovery and formalization of the phenomenon of control is the first alternative to the cause-effect perspective ever proposed in any science.

The first discussion of purposeful machines and people came in 1943 in a paper called: Behavior, Purpose and Teleology by Rosenblueth, Wiener and Bigelow. This paper also argued that purpose belongs in science as a real phenomenon in the present. Purpose does not mean that somehow the future influences the present.

The first specific suggestion on how to use the concept of control to understand people came in 1957 in a paper entitled: A General Feedback Theory of Human Behavior by McFarland, Powers and Clark.

In 1973 William T. (Bill) Powers published a seminal book called "Behavior: the Control of Perception," which still is the major reference for PCT. See literature below.

This book spells out a complete model of how the human brain and nervous system works like a living perceptual control system. Our brain can be viewed as a system that controls its own perceptions. This view suggests explanations for many previously mysterious aspects of how people interact with their world.

Perceptual Control Theory has been accepted by a few independently thinking psychologists, scientists and other interested people. The result is that an association has been formed (the Control System Group), several books published, this CSGnet set up and that at latest count 16 professors are teaching PCT in American universities today.

DEMONSTRATING THE PHENOMENON OF CONTROL

The phenomenon of control is largely unrecognized in science today. It is not well understood in important aspects even by many control engineers. Yet the phenomenon of control, when it is recognized and understood, provides a powerful enhancement to scientific perspectives.

It is essential to recognize that this phenomenon exists and deserves an explanation before any of the discourse on CSGnet will make sense.

Please download the introductory demonstration demla.exe, which is an interactive program and/or rubberbd.txt, which is a text telling you how to demonstrate the phenomenon to yourself and a friend using only two rubber bands.

THE PURPOSE OF CSGnet

CSGnet provides a forum for development, applications and testing of PCT.

CSGnet PARTICIPANTS

Many interests and backgrounds are represented here. Psychology, Sociology, Linguistics, Artificial Intelligence, Robotics, Social Work, Neurology, Modeling and Testing. All are represented and discussed. A quality of participants on this net is that most are prepared to question and reconsider what they think they know.

HOW TO ASK EFFECTIVE QUESTIONS

Since PCT puts much conventional, well established wisdom on its head, it is helpful to begin by demonstrating the phenomenon of control to yourself and studying a few references. As you catch on, read this net and follow a thread that interests you for a month or more.

When you ask a question, please consider that in order to give you a good answer, a respondent will need to put your question in context.

Therefore, please introduce yourself with a statement of your professional interests and background. It will help if you spell out which demonstrations, introductory papers and references you have taken the time to digest and what you learned.

People on this net are in various stages of learning and understanding PCT. When you get a reply to your post, please consider that the respondent who found your question of interest and invested time in a reply, may benefit from knowing how you perceived the answer. Did it answer your question? Was it clear? Were you able to understand it?

THE CONTROL SYSTEMS GROUP

Serious enthusiasts of PCT have formed the Control Systems Group. Annual dues are \$45 for regular members and \$5 for students (subsidized). This group meets once a year (1993: July 28-Aug 1) in Durango, Colorado, for informal presentations and exchanges. The group also publishes threads from this net on a quarterly basis in a booklet called the Closed Loop. These booklets are distributed to members and are available separately for a fee. For membership information and back issues of Closed Loop, write: CSG, c/o Mary Powers, 73 Ridge Place CR 510, Durango, CO 81301-8136.

SUBSCRIBING TO CSGnet

When you subscribe to CSGnet, you get this message, CSGINTRO.DOC. But you may have received it from a friend who printed it, seen it on a demodisk, or seen it on USEnet. To subscribe, send a message as follows:

To: (Internet) LISTSERV@VMD.CSO.UIUC.EDU
Subscribe CSG-L Lastname, Firstname, Affiliation, City, State
help

(Technically, your last name, first name etc is optional). ("help" requests a list of most commonly used commands). (The Bitnet address is: LISYSERV@UIUCVMD). (This server is not sensitive to upper or lower case letters).

CSGnet can also be accessed via Usenet (NetNews) where it is listed as the newsgroup bit.listserv.csg-1.

HOW TO OBTAIN TEXT AND PROGRAM FILES

A number of ASCII documents and binary computer programs are available on a fileserv maintained by Bill Silvert. It is possible to download all these files via e-mail. If you are on internet, it is easiest to obtain binary program files via anonymous FTP. If you are on MCI mail, you have read about how you can transfer binary files with Kermit or Zmodem protocols. (Type help at the MCI mail prompt for directions). But the server cannot send binary files over the internet mail network, so download uue.scr first, then request the binary files uuencoded as ASCII files. The Internet address for the server is BIOME.BIO.NS.CA. CSGnet files are kept in the subdirectory pub/csg.

To get basic information and a current listing of available documents, send a message as follows:

To: (Internet) SERVER@BIOME.BIO.NS.CA.
help

get csg/Index

"help" asks the server to send you commands and explanations.
"ftp" requests the scoop on anonymous FTP for internet.
"get csg/Index" requests the Index for the csg subdirectory.

In your message, pay attention to upper and lower case! DOS is not dos.

As part of the index (of the csg directory), you may be looking at:

programs/msdos: demla.exe 128437 Bill Power's demonstration of
perceptual control dem2a.exe 123649 Bill Power's modeling of control

programs/source: uud.scr 54371 ASCII. Compile uud.exe w/ DOS
debug. Directions @ end.

If you want to request demla.exe (uuencoded) to get a "live" demonstration of the phenomenon of control, and the ASCII file uud.scr which allows you to use DOS debug to create uud.exe to decode it, send the following message:

uue csg/programs/msdos/demla.exe get csg/programs/source/uud.scr

The uuencoded demla.exe will be sent in four parts. Remove headers and use

an editor to make it into one file (starting with table and ending with end) before you use uud.exe to restore the file. demla.exe is a self-extracting archive file. Put it in it's own directory before you execute it. You get complete documentation and a running program.

LITERATURE REFERENCES

For a complete list of CSG-related publications, get the file biblio.pct from the fileserver as described above. Here are some selected books on perceptual control theory:

Powers, William T., Behavior: The Control of Perception. Hawthorne, NY: Aldine DeGruyter, 1973, 296 pages. The foundation of PCT! A seminal book.

Robertson, Richard J. and Powers, William T., editors. INTRODUCTION TO MODERN PSYCHOLOGY; The Control Theory View. CSG Books (\$25 postpaid) 1990, 238 pages. Textbook on psychology for universities. Highly recommended.

Powers, William T. LIVING CONTROL SYSTEMS I: Selected Papers. CSG Books 1989, (\$16.50 postpaid) 300 pages. A collection of previously published papers.

Powers, William T. ,LIVING CONTROL SYSTEMS II: Selected Papers. CSG Books (\$22 postpaid) 1992, 274 pages. A collection of previously unpublished papers.

Marken, Richard S, MIND READINGS: Experimental Studies of Purpose. CSG Books (\$18 postpaid) 1992, 212 pages. Research papers exploring control.

Marken, Richard S., editor. PURPOSEFUL BEHAVIOR: The Control Theory approach. Thousand Oaks, CA: Sage Publications: American Behavioral Scientist, special issue. Vol. 34, Number 1. September/October 1990. 11 articles, 16 contributors, 121 pages. A very readable introduction to a science of purpose and supportive research. Highly recommended.

Runkel, Philip J., CASTING NETS AND TESTING SPECIMENS. New York: Praeger, 1990, 186 pages. Contrasting the proper and improper uses of statistics with modeling for understanding and prediction of people as well as processes. Highly relevant to TQM efforts!

Hershberger, Wayne, editor, VOLITIONAL ACTION, CONATION AND CONTROL. Advances in Psychology 62. NY: North-Holland, 1989. 25 chapters, 33 contributors, 572 pages.

Ford, Edward E., FREEDOM FROM STRESS. Scottsdale AZ: Brandt Publishing, 1989, 184 pages. A highly readable introduction to PCT and a personal problem solving guide. The most accessible text available. Written as a comprehensive counseling story anyone can relate to. Most recommended!

Gibbons, Hugh, THE DEATH OF JEFFREY STAPLETON; Exploring the Way Lawyers Think. Concord NH: Franklin Pierce Law Center, 1990, 197 pages. Textbook for law students which spells out how lawyers think by explaining and using a PCT framework.

McClelland, Kent, PERCEPTUAL CONTROL AND SOCIOLOGICAL THEORY. 1992. This

unpublished paper suggests that individual control as a phenomenon is central to understanding sociology. Available from the author.

McPhail, Clark, THE MYTH OF THE MADDING CROWD. Hawthorne, NY: Aldine de Gruyter, 1990, 265 pages. Explains group behavior as a function of purposeful individuals.

Petrie, Hugh G., DILEMMA OF ENQUIRY AND LEARNING. Univ. of Chicago press, 1981. Discusses learning with explicit recognition of PCT insight.

Richardson, George P., FEEDBACK THOUGHT IN SOCIAL SCIENCE AND SYSTEMS THEORY. Univ. of Pennsylvania Press, 1991, 374 pages. A review of systems thinking in history, cybernetics, servo mechanisms and social sciences. Provides a perspective placing PCT in context in relation to other paradigms of human behavior.

CSG Book publishing: 460 Black Lick Road, Gravel Switch, KY, 40328

Date: Wed Nov 11, 1992 8:18 am PST
Subject: associative memory

[From: Bruce Nevin (Wed 921111 10:12:08)] (Martin Taylor 921110 14:30) --

In re "meaningful or less" (Bruce Nevin 921110 12:55:31)

I've been "leaning on" something called associative memory for some time in my efforts to understand the relation between nonverbal perceptions and words. In Bill's proposal, a word is associated with nonverbal perceptions in the input function of a category-level ECS. This seems to me more crisp and well defined than is suggested by my subjective experience of associative memory. Perhaps I am underestimating the contribution of imagination, and the amplification (or cascading, rather) of that contribution through the interconnectedness of ECSs at the category level. Whatever the case may actually turn out to be, I have assumed that associative memory introduces (and functions in) pathways over and above those used by the familiar p, r, and e signals in our diagrams of an ECS.

This then is the answer to your question:

>> In PCT, meanings are remembered or imagined
>>perceptions associated with a given perception.
>
>Sorry? What construct in PCT is an association? This sentence doesn't fit
>my simple view of PCT as an acyclic network of ECSs that deal in the control
>of scalar perceptual signals.

Bruce bn@bbn.com

Date: Wed Nov 11, 1992 8:32 am PST
Subject: Re: Primer for modelers: draft

POWERS_W%FLC@VAXF.COLORADO.EDU (William T. Powers) writes:

```
:  
:  
:           Experimenting with the control paradigm  
:           A primer for computer modelers  
:           DRAFT: William T. Powers  
:  
: [ beginning part deleted ...]  
  
: The simplest control system to model on a digital computer is one in  
: which all the functions are simple proportionalities except the  
: output function, which is an integrator. Alternatively, the feedback  
: function or the input function can be made into an integrator;  
: however, only one function should be an integrator and the rest  
: should be proportional multipliers. We will use a design with an  
: integrator in the output function; you can experiment with the other  
: possibilities.
```

Why should there be only one integrator?

..stuff deleted..

```
: NOTE: I would appreciate it if programmers versed in BASIC, Pascal,  
: and other languages would write versions of the above program, test  
: them, and transmit them to me for inclusion in an appendix to this  
: primer. Programs should be as generic as possible so as to run on as  
: many computers as possible. I especially need versions for mainframes  
: and workstations. In the future I will simplify the program and make  
: it easier for the user to change parameters from the keyboard. For  
: now, I hope that all programmers and would-be programmers will try  
: the program as it stands and learn from it, even if you have to get  
: some help from a 14-year-old. As this primer evolves I will post  
: revisions and again ask for your help.  
: -----
```

Here is a PASCAL-Version of your program. It should work on whatever platform ..

```
program control;  
const kd = 1.0;  
      ki = 1.0;  
      ko = 8.0;  
      kf = 1.0;  
      dt = 0.1;  
  
var   sp, sr, se : real;  
      qi, qo, qd : real;  
      i : integer;  
  
begin  
  (* initialize variables *)  
  sp := 0.0;  
  sr := 20.0;  
  se := 0.0;
```

```
qo := 0.0;
qi := 0.0;
qd := 0.0;

for i := 0 to 24 do begin
  if i > 12 then qd := 10.0 else qd := 0.0;
  qi := kf*qo + kd*qd;
  sp := ki*qi;
  se := sr - sp;
  qo := qo + ko*se*dt;

  write ('qd = ',qd:6:2,' qi = ',qi:6:2,' qo = ',qo:6:2);
  writeln (' sp = ',sp:6:2,' sr = ',sr:6:2,' se = ',se:6:2);
end;
readln;
end.
```

```
: Best,
:
: Bill P.
```

--

Wolfgang Zocher
Regionales Rechenzentrum f. Niedersachsen
Schlosswender Strasse 5
3000 Hannover

Date: Wed Nov 11, 1992 9:42 am PST
Subject: unwanted meanings

[From: Bruce Nevin (Wed 921111 10:37:20)] (Bill Powers (921110.1430) --

```
>Do you agree that we also control for unwanted meanings?
>I.e., some traversals lead to words that have unwanted meanings that
>stick out enough to cause us to switch to a different traversal, or to
>insert a specific denial of the unwanted meaning, or to loop back and
>repeat part of the traversal in different terms.
```

Yes, skill in using language is largely avoiding such built-in pitfalls.

```
>>It seems to me that we use the socially learned structure of
>>language as an framework for organizing how our attention traverses
>>our perceptions (the combination of perceptions from the
>>environment with perceptions from memory and imagination).
>
```

```
>Yes, I agree that when we think in words (not always the case) we use
>the socially (well, individually) learned structure as a way of guiding
>the traversal of meanings -- for creating a scenario to go with the words
```

I suspect that reference signals that get set when we tell stories about our perceptions persist when we're not thinking in words.

```
>>Failure of social agreements is probably occasion for intrinsic >error
>in mammals (possibly in other creatures as well).
```

>

>I can think of lots of ways in which such failure would lead to
>intrinsic error without the failure itself being an intrinsic error.

Bateson was starting a lecture and suddenly began demanding of the (small) audience "What are you sitting there looking at me for?" etc. in a way that they did feel (as they afterwards confessed) that they were doing something inappropriate. He then went on to talk about this distress at being wrong--existentially wrong, not belonging, odd man out goes y o u yes I mean YOU! This feeling of distress, which I suspect we can all identify in our memories somewhere or other, does not so far as I can tell depend upon consequences of a failure of agreement nor even upon such consequences as anticipated in imagination. It is very immediate. Many social arrangements, from club membership to the cultivated persona of a tough, independent outsider, are strategies for avoiding this distress. Peer pressure, and all that.

Bruce bn@bbn.com

PS--oh, yes, Bateson talks about this in connection with mammalian communication, not just human double bind, etc.

Date: Wed Nov 11, 1992 10:18 am PST
Subject: Introducing myself

(gabriel 921111 0945CST)

I realised last night that I had started a conversation on the net without proper introduction. Tsk tsk.

I am a computer jock, close to retirement. Bill Cunningham introduced me to the net. In the past I have made my living as an Electrical Engineer and as a theoretical physicist doing Group Theory and Quantum Mechanics.

These days I work on software questions of various kinds. The proximate cause of my participation in CSG is an interest in decision support, and, more remotely, a concern to have wise government, so as to avoid civil unrest, war, and other things we would all sooner not have happen. It seems to me that fits well with the etymology of "Cybernetics."

I met Bill C. some years ago when we both laboured to see that Uncle Sam's money was wisely spent in a project whose intention was to extract needles from haystacks of messages about battles. Not very different from extracting needles of truth from haystacks of press releases about elections.

Now, I work on proofs about computer programs for amusement - hence the interest in attribute spaces and Hamming distances, and in information theoretic aspects of the theory of computation. From there it is not a great leap to PCT, and to application of ideas from PCT in decision support.

John Gabriel (gabriel@eid.anl.gov)

Date: Wed Nov 11, 1992 11:29 am PST
Subject: o tell me where is meaning bred

[From: Bruce Nevin (Wed 921111 10:59:45)] (penni sibun 921110.1600) --

> the concept of ``meaning,' ' or anything else, is useful in a theory
> only insofar as it has explanatory power. you seem to be saying here
> that meaning is everything--is inherent. to say that something is
> everything isn't much different from saying it is nothing: what good
> does it do to say ``that has property x' ' when everything has property
> x?

No, I am not saying that meaning is everything, I am saying that each thing "has" specific meanings. Understand that every "thing" is a perception, and each meaning is another perception. I am not saying that a given thing/perception "has property x" i.e. has meaning in some vague, generalized sense of "the property of meaning." I am saying that each given thing-perception has associated with it specific other thing/perceptions, differently for each, and not verifiably the same in each person. I believe that this treatment does have explanatory power, and is very different from denying the existence of meaning, as you seem to be doing. Furthermore, since verbal perceptions are distinguished in important ways from nonverbal perceptions (see below), the meanings that can be associated with a given language-perception are organized (by the structure in language) in ways that the meanings that can be associated with a given nonverbal perception cannot be.

A word W "means" certain other perceptions for a given person at a given time.

Word W: all the perceptual and reference signals involved in the person saying W, and all of the signals involved in the person recognizing W.

other perceptions: perceptual and reference signals associated with the word W in that person (by associative memory, in input functions, by some other TBD means).

That person and others are most likely able to come to agreement about such of those perceptions (separately in each of them) as are involved in their cooperative action. That is, they are likely to agree as to what W means with respect to their cooperative action. They are likely to agree even when they think they detect some discrepancies, for the sake of carrying out those cooperative activities. (If on the other hand there is some reference signal for failure of cooperation, such small differences are likely to be perceived as large and important-- what Freud called the Narcissism of small differences, something we all know about. But let's keep it simple.) In this process, they come to perceive the meanings as being externalized from themselves, either as natural facts or as social facts. In every instance of cooperative action people perceive themselves as depending upon this seemingly external framework of meanings. In fact, during the course of each instance of cooperative action the participants each individually and privately test and elaborate their meaning-perceptions, the very ones which they had presumed a pre-existent external reality to which they must conform.

Here's the kicker, Bill: insofar as all the other participants are doing the same, they are correct in this presumption.

> it's not clear to me whether you think
> speech actions are distinguished actions or not. (my argument is that
> they are not distinguished (in any nonincidental way)).

In general, I am not concerned with actions whether verbal or nonverbal, but with perceptions, the control of which has actions as a byproduct. Verbal perceptions are distinguished from most nonverbal perceptions in that they participate in a conventionalized structure by which people create and transmit information. There are iconographic systems that are not so articulate, and there are language-like artifacts (notably mathematics and logic) that lack the flexibility and informational capacity of language, which are also distinct from other nonverbal perceptions insofar as they are structured according to agreed conventions. References on request.

I am sure that sometime in the future I will be able to put all this more clearly and succinctly.

Bruce bn@bbn.com

Date: Wed Nov 11, 1992 11:36 am PST
Subject: Meaning; Second Order Cybernetics

[From Bill Powers (921111.0900)]

On meaning:

As I see it from the HPCT standpoint, meaning is not something objective; words don't "have" meanings. Meanings are the perceptions evoked in imagination by the symbol that is supposed to have a meaning.

Avery Andrews (921110):

>Consider the meaning of the NP `the Iliad'. There are many different
>kinds of perceptions that this can be hooked up to. In the benign
>environment of your local library, you can ascertain that a book is the
>Iliad by noting that it has `Iliad' written on its spine, or by opening
>up and seeing that the text starts out `Sing, Muse, of the wrath of
>Achilles', So a wide range of different perceptual functions will
>suffice distinguish copies of the Iliad from other things ...

If you think that the book you are holding is called 'the Iliad,' then that is the meaning of "Iliad" for you: it means the book you are holding and see in front of you. The word also means whatever else you are reminded of, which depends on your experience and memories.

>On the other hand, you can imagine more treacherous worlds, where
>books say on their covers that they are the Iliad, but have something
>completely different inside, or have all sorts of omissions and
>interpolations in the text, so that some expertise (e.g. a large and
>sophisticated perceptual function) is required to distinguish good
>copies of the Iliad from bad ones (the Ancient world was like this,
>before the Alexandrians cleaned up the text).

But that objection is based on the idea that there is a "real" meaning of "The Iliad" that can be objectively established by experts. All this shows is that different people will experience different meanings of the same term, and if some of these people agree on the meanings, and can persuade others to consider them as authorities, they may be able to influence others to adopt (some of) the same meanings. Even so, if you hold a book in your hand that says "The Iliad" on the cover, the meaning of "Iliad" will be what you're

experiencing -- so far. If you're an expert, you might add under your breath "_purported_ Iliad." If so, then the book you're looking at is the meaning of "_purported_ Iliad." Meaning is perception. What else could it be?

Penni Sibun (921110.1600) --

>the concept of ``meaning,' ' or anything else, is useful in a theory
>only insofar as it has explanatory power.

But meaning is not a theoretical concept -- it's a phenomenon. When I say words to you, you experience more than the words. You experience their referents -- their meanings. If you don't experience any meanings of words, then you don't understand them; they might as well be in a foreign language (you being a linguist, I must add " -- that you don't know"). Control theory provides an explanation of the phenomenon of meaning, not the other way around. When I say "Penni's father" to you, you know what those words mean -- they call to your mind a specific person who looks and acts in familiar ways. That is the phenomenon to be explained. Theories are supposed to explain phenomena. I believe that we have at least the beginnings of a control-theoretic explanation of the phenomenon of meaning.

RE: Carver & Scheier and Vallacher: all of these people started their writing about control theory after learning about my work and discussing it with me. They understand it to a certain degree, but are trying to merge it into existing methodologies and concepts. This has never succeeded, and they do not succeed, either. I have invited Carver & Scheier to meetings of the CSG, but they have been too busy to come.

RE: Second order cybernetics etc.

Margaret Mead suggested second-order cybernetics -- applying cybernetic concepts to cyberneticists, or "Cybernetics of Cybernetics" -- at the first annual symposium of the American Society for Cybernetics. The proceedings, published in 1968, nowhere mention the date or location of this symposium, but it must have occurred before 1968. They appear in Purposive Systems, edited by von Foerster, White, Peterson, and Russell, and published by Spartan Books (1968). At this time my work on what is now called PCT had been in progress for 15 years, but it is not mentioned.

In this field Mead was a dilettante as were many of the others at the symposium. She was just trying to get one-up on all these cyberneticists, whom she didn't understand at all. It was in this volume (p. 26) that the neurologist Ralph W. Gerard said "I have always regarded a drop of water sliding down a slightly inclined plane as showing all the manifestations of purposeful behavior." "Always," I presume, meant both before and after learning about cybernetics, sufficient proof to me that Gerard knew nothing about purpose and what cybernetics had to say about it (or might have had to say). At some time or other, Gerard became a President of the ASC as a reward, I suppose, for his ignorance.

Cyberneticists have been floating upward into the cloudy realms of philosophy ever since, when they have not been promoting a technocracy or supporting Marxist revolutions by implementing expensive and unworkable central computer control of the economy. Of control theory they would know essentially nothing, were it not for Cliff Joslyn's recent work. I spent many years attending cybernetics meetings and even giving the occasional paper, but my message was generally received with hostility or simply ignored. Only a few people, like Heinz von Foerster, gave me any support. All the support faded away when it was realized that I did not buy Ashby's conclusions about control theory, that I thought Maturana was confused, and that I thought control theory was more important than recursive self-computations of awareness.

For the most part there is nothing of interest to me going on in cybernetics.

Best to all, Bill P.

Date: Wed Nov 11, 1992 11:37 am PST
Subject: meaning, meaning, who's got the meaning?

[From: Bruce Nevin (Wed 921111 11:48:22)]

(Avery Andrews (Wed, 11 Nov 1992 13:22:28 EST)) --

So meaning depends on perceptions, but not necessarily perceptions of any of the people actually involved in a linguistic communication where the meaning figures.

This is "meaning" understood from the perspective of a supposed external observer. I believe that as we each create meanings we presume that they are external, universal properties which we are (imperfectly) learning.

Per the "kicker" in my aside to Bill, I believe that in our cooperative activities we are in fact collectively creating commonalities of meaning, though there is no way of being certain of that, and the process never reaches closure, so it is not imperfect learning but rather never completed creation. This I think is close to Sapir's perspective, refs and quotes if you want (but not right away).

Bruce bn@bbn.com

Date: Wed Nov 11, 1992 3:04 pm PST
Subject: primer program from Zocher

[From Bill Powers (921111.1230)] Wolfgang Zocher (921111) --

RE: primer for computer modelers

Thanks for the Pascal version, Wolfgang. I will incorporate it, with credit. It runs fine under Turbo Pascal -- anybody else out there who can test these programs to make sure they work on your system?

I would welcome the efforts of anyone who would like to take my rather verbose version of the writeup and make it shorter and better organized.

Also, if someone wants to make the program more user-friendly (but still short), feel free: I'll check it out and revise the Primer accordingly. I can run C, Pascal, and Qbasic programs.

You ask why there shouldn't be a second integrator. This should go into the Primer. The best way to see why not is to put a second integrator into the program.

Try this: in place of

qi := kf*qo + kd*qd;

introduce a dummy variable x, and write

```
x := x + kf*qo;  
qi := x + kf*qd;
```

You'll have to declare x and initialize it to 0. The variable x represents the contribution of the output via the feedback function to the state of the input quantity. The feedback function is now an integrator.

How about trying this, and writing back to tell the folks what happened?

(C version: same change, but use = instead of :=).

Best, Bill P.

Date: Wed Nov 11, 1992 3:46 pm PST
Subject: MESSAGE FROM MARY

[from Mary Powers 921111] Jixuan Hu:

You can tell that Bill Powers is pretty disenchanted with the ASC. Actually, the Control Systems Group was, in 1983 and 1984, a group within the ASC, but after two meetings in which we found we were only talking to each other, we began to meet independently. Some of us attended further ASC meetings, and you might want to look up The Conference Workbook for "Texts in Cybernetics" for the ASC meeting in Felton, California if you want to compare Glassersfeld, Maturana, and Powers. The Powers text from that book is also reprinted in Living Control Systems I. But you should read his Behavior: the control of perception.

The ASC deals in abstract generalizations, and PCT is concerned with modelling organization. They both are founded on Wiener's book, which by the way is about control and communication, not communication and control. But cybernetics has always downplayed control, while PCT considers it fundamental.

Everybody:

I am getting requests for information about CSG from netters and by snail mail. The following is what I am sending out (the bibliography is shamelessly cribbed and edited from Dag's). I would appreciate suggestions and revisions (keeping in mind that the PCT section is at max length since I have postage considerations and can fit everything now on two double-sided pages)

Perceptual Control Theory

While the existence of control mechanisms and processes (such as feedback) in living systems is generally recognized, the implications of control organization go far beyond what is generally accepted. We believe that a fundamental characteristic of organisms is their ability to control; that they are, in fact, living control systems. To distinguish this approach from others using some version of control theory but forcing it to fit conventional approaches, we call ours Perceptual Control Theory, or PCT.

PCT requires a major shift in thinking from the traditional approach: that what is controlled is not behavior, but perception. Modelling behavior as a dependent variable,

as a response to stimuli, provides no explanation for the phenomenon of achieving consistent ends through varying means, and requires an extensive use of statistics to achieve modest (to the point of meaningless) correlations. Attempts to model behavior as planned and computed output can be demonstrated to require levels of precise calculation that are unobtainable in a physical system, and impossible in a real environment that is changing from one moment to the next. The PCT model views behavior as the means by which a perceived state of affairs is brought to and maintained at a reference state. This approach provides a physically plausible explanation for the consistency of outcomes and the variability of means.

The PCT model has been used to simulate phenomena as diverse as bacterial chemotaxis, tracking a target, and behavior in crowds. In its elaborated form, a hierarchy of perceptual control systems (HPCT), it has lent itself to a computer simulation of tracking, including learning to track, and to new approaches to education, management, and psychotherapy.

Control systems are not new in the life sciences. However, numerous misapprehensions exist, passed down from what was learned about control theory by non-engineers 40 or 50 years ago without further reference to newer developments or correction of initial misunderstandings. References in the literature to the desirability of positive feedback and the assertion that systems with feedback are slower than S-R systems are simply false, and concerns about stability are unfounded.

The primary barrier to the adoption of PCT concepts is the belief - or hope - that control theory can simply be absorbed into the mainstream life sciences without disturbing the status quo. It is very hard to believe that one's training and life work, and that of one's mentors, and their mentors, must be fundamentally revised. Therefore, PCT appeals to those who feel some dissatisfaction with the status quo, or who are attracted to the idea of a generative model with broad application throughout the life sciences (plus AI and robotics). There are very few people working in PCT research. Much of its promise is still simply promise, and it meets resistance from all sides. It is frustrating but also tremendously exciting to be a part of the group who believe that they are participating in the birth of a true science of life.

BIBLIOGRAPHY

Powers, William T., Behavior: The Control of Perception.
Hawthorne, NY: Aldine DeGruyter, 1973.
The basic text.

*Robertson, Richard J. and Powers, William T., editors.
Introduction to Modern Psychology: the Control Theory View. CSG
Books, 1990. (\$25 postpaid)
College-level text.

*Powers, William T., Living Control Systems I : selected papers.
CSG Books, 1989. (\$16.50 postpaid)
Previously published papers, 1960-1988.

*Powers, William T., Living Control Systems II : selected papers.
CSG Books, 1992. (\$22 postpaid)
Previously unpublished papers, 1959-1990.

*Marken, Richard S, Mind Readings : experimental studies of purpose. CSG Books, 1992. (\$18 postpaid)
Research papers exploring control.

Marken, Richard S., editor. Purposeful Behavior : the control theory approach. Thousand Oaks, CA: Sage Publications: American Behavioral Scientist, special issue. Vol. 34, Number 1. September/October 1990.
11 articles on control theory.

Runkel, Philip J., Casting Nets and Testing Specimens. New York: Praeger, 1990., 186 pages.
When statistics are appropriate; when models are required.

Hershberger, Wayne, editor, Volitional Action : conation and control. Advances in Psychology 62. NY: North-Holland, 1989.
25 articles (not all PCT)

Ford, Edward E., Freedom From Stress. Scottsdale AZ: Brandt Publishing, 1989.
A self-help book. PCT in a counseling framework.

Gibbons, Hugh, The death of Jeffrey Stapleton : exploring the way lawyers think. Concord, NH: Franklin Pierce Law Center, 1990.
A text for law students using control theory.

McClelland, Kent, Perceptual Control and Sociological Theory. 1992.
Not yet published. Available from the author, Grinnell University, Grinnell, Iowa.

McPhail, Clark, The Myth of the Madding Crowd. Hawthorne, NY: Aldine de Gruyter, 1990.
Introduces control theory as a way to explain group behavior.

McPhail, Clark, Powers, William T., and Tucker, Charles W., Simulating Individual and Collective Action In Temporary Gatherings. Social Science Computer Review, 10:1, Spring 1992. Durham, NC, Duke University, 1992.
Computer simulation of control systems in groups.

Petrie, Hugh G., Dilemma of Inquiry and Learning. Univ. of Chicago press, 1981.
Introduces PCT to educational theory.

Richardson, George P., Feedback Thought in Social Science and Systems Theory. Univ. of Pennsylvania Press, 1991.
A review of systems thinking, including PCT.

*These items are available from CSG Books, 460 Black Lick Road, Gravel Switch, KY, 40328.

The CSG is an organization of people in the behavioral, social, and life sciences who see the potential in PCT for increased understanding in their own fields and for the unification of diverse and fragmented specialties. We maintain a computer network, and also have a publication, Closed Loop: threads from CSGnet. Net subscribers find it useful to have thematic collections of some of the network discussions, and it enables non-net members to keep up with them.

Closed Loop is published quarterly. It is sent only to members, who pay \$45 a year. This sum supports student subscriptions, which are \$5. Any surplus after production and mailing, plus a few organizational expenses, is used to help defray costs for students wishing to attend the annual meeting.

The annual meeting is held in Durango, Colorado, on the campus of Fort Lewis College. In 1993 it will begin in the evening of Wednesday, July 28, and end Sunday morning, Aug. 1. There will be 7 plenary meetings (mornings and evenings), with afternoons, mealtimes, and late night free for further discussion or recreation. Full details will be available on the net or by mail after April 1, 1993.

To subscribe to the net, this is the address:

LISTSERV@VMD.CSO.UIUC.EDU

send this message:

SUBSCRIBE CSG-L Lastname Firstname Location/Affiliation

To join CSG, receive conference information, etc., write to

The Control Systems Group
Mary A. Powers
73 Ridge Place
Durango CO 81301-8136

A complimentary copy of Closed Loop will be sent upon request. Back issues are available: Volume 1 (4 issues) is \$12. Single issues of Volume 2, beginning with Jan. 1992, are \$6 each.

penni:

I appreciate your gripe about the cost of CSG membership and hereby offer you a limited-income membership for \$20. (Anyone who objects is welcome to impeach me. Chuck? Wayne?)

Rick:

>You can't tell what people are doing by watching what they are doing.

No indeedy. You apply the Test for the Controlled Quantity. I ran across a nice quote about the Test. Almost 3 centuries old.

Search then the ruling passion; there alone
The wild are constant, and the cunning known;
The fool consistent, and the false sincere;

Priests, princes, women, no dissemblers here.
This clue once found unravels all the rest.

Alexander Pope, 1688-1744

Happy unravelling! Mary P.

Date: Wed Nov 11, 1992 6:56 pm PST
Subject: meaning

[Avery Andrews 921112.1238] (Bill Powers (921111.0900))

>If you think that the book you are holding is called 'the Iliad,' then
>that is the meaning of "Iliad" for you: it means the book you are
>holding and see in front of you. The word also means whatever else you
>are reminded of, which depends on your experience and memories.

True. But there is another rather interesting sense of `meaning' in which your meaning is *wrong*. E.g. if you used it as a basis for teaching in a Classics department, it would be quite in order for you to be sacked for incompetence.

>But that objection is based on the idea that there is a "real" meaning
>of "The Iliad" that can be objectively established by experts. All this

Which there is, sort of, tho `objective' might not be quite the right word for. The real Iliad is whatever Homer dictated to however wrote it down, on current conceptions. Being a bit more precise, people who are into this sort of stuff seem to agree (at least implicitly) that if somebody composed something pretty much like the Iliad as we have it and dictated it to a scribe, or wrote it himself, or got it to be memorized & transmitted by a school of apprentices for a while before it got written down (the first possibility is presently considered the most likely, as far as I can make out), then that person is Homer & what he composed is the real Iliad. If none of these presumptions are true, then maybe there isn't a real Iliad.

I'm not denying that all this is ultimately based on perception, but I'm trying to emphasize a social aspect of meaning that PCT maybe puts a bit too much in the background.

Avery.Andrews@anu.edu.au

Date: Thu Nov 12, 1992 4:58 am PST
Subject: consensual reality

[From: Bruce Nevin (Thu 921112 07:25:24)]

(Avery Andrews 921112.1238)

> there is another rather interesting sense of `meaning' in
>which your meaning is *wrong*. E.g. if you used it as a basis for
>teaching in a Classics department, it would be quite in order for you
>to be sacked for incompetence.

Demonstrating competence as an employee is a special case of demonstrating membership. One must demonstrate to others (those who matter) that one controls certain perceptions to their satisfaction.

Ah yes: one controls one's own perceptions so that those others experience no error in their own perceptions. One controls one's perception of their experiencing no error . . . Lots of room here for illusion, delusion, projection, reification. But I agree: there does seem to be some supra-individual collective creation that persists over time and is shared across persons. Sapir's view I think (though he obscures things when he waxes poetic, as he likes to do), and Harris's, is that such entities as "the grammar" of "the language," "the culture," and indeed "the Iliad," are artifacts of the science or other study that investigates them. Investigators into (their private) perceptions of some particular thing negotiate agreements about that thing, and then reify their negotiated agreements as the thing itself. The consensus- perceptions are felt to be different from the private perceptions, and indeed they are unless one happens to agree with the consensus utterly. In particular, negotiated consensus is much more stable than are ongoing subjective perceptions. This stability and other sensed differences are attributed to the thing that is perceived as apart from one's subjective perceptions of that thing; but the consensus is itself a perception, and not the Ding an sich.

>I'm not denying that all this is ultimately based on perception, but I'm
>trying to emphasize a social aspect of meaning that PCT maybe puts a bit
>too much in the background.

What I have been trying to say about this: while we imagine that we are learning it, we create it. In the course of acting in concert with others, while imagining that we or others had failed to learn it adequately, we test, renegotiate, and refine it.

(BTW, is Lord's work on oral-traditional bardic performance in Yugoslavia and Macedonia now passe'?)

Bruce bn@bbn.com

Date: Thu Nov 12, 1992 5:43 am PST
Subject: Re: tutorial on modelling

[From: Bruce Nevin (Thu 921112 08:10:20)]

Bill,

Thank you for that excellent tutorial on modelling. It will be a vade mecum reference document for me for some time. Maybe I'll get a compiler going sometime soon.

Especially helpful was seeing how each function can know of its output only by way of its input, which carries the ongoing cumulative effect of the entire loop. The input looks like its back door to us, as external observers, who are interested in the outputs of the function; but the function can be "interested" only in its input, since it is to its input that it responds. In this, a function in an ECS is analogous to a control system: we are mistaken when we, as external observers, are surprised that the effect of a change comes round the loop to its "back door." Its perceptual input is really its "front door." Both the control system and a function within it face upstream.

When we're dealing with inanimate objects, we are rightly concerned with their outputs (actions, behaviors, etc.). When we're dealing with living things, we should concern

ourselves with their inputs, and with divining their controlled inputs. Alice's venture through the looking- glass is nothing to this. The reversal is existentially shocking. But the shock is I think intellectual. "Intuitively" I suspect we know how to do this; we're control systems, after all. So I think a good modus operandi is to observe (without interference) how we do when we are doing well with respect to living beings. It then becomes crucial to distinguish the social conventions by which we continually grease the skids. So I mean we need to observe how we do well in those gaps not covered by conventional interaction.

Bruce bn@bbn.com

Date: Thu Nov 12, 1992 6:45 am PST
Subject: ends-based public policy

[From: Bruce Nevin (Thu 921112 09:15:48)]

This is from a forum on education reform (kl2@world.std.com). The writer asks for advice on building community support and implementing. How well does it measure up as an instance of the kind of ends-directed public policy that Rick was advocating re Clinton? Are these perceptual outcomes about which people can productively come to agreement? Are they controllable by the parties involved (teachers, students, parents, administrators) through their cooperative or noncooperative action? If not, why not, and what would improve the approach?

--+=-----+=-----+=-----+=-----+=-----+=-----+=-----+=-----+=-----+=-----

. . .

For those . . . not familiar with outcomes-based education, let me give some background. Outcomes are a set of characteristics that all students are expected to acquire through the educational process, including learning going on in the home and outside the school. My understanding of this approach is that the set of outcomes provides a coherent framework on which curriculum and program decisions can be made. In essence, it is an educational philosophy statement for the school. In making choices regarding what to teach, appropriate personnel can look to the outcomes statement for guidance. At least, that is my interpretation, but I hasten to add that I am not formally trained in education. . . . Here are the Millis schools outcomes, still in draft form:

Millis Superintendent's Citizens Advisory Council

Draft 3 of Student Outcomes 7/30/92

The Millis Superintendent's Advisory Council, a broadly-based group, with representatives from the schools and the community, was formed in September of 1991. This group worked towards defining a set of outcomes appropriate for students of the Millis School System. These outcomes define the set of characteristics we expect all students to acquire through the educational process.

In order to achieve these outcomes, both teachers and administrators are charged with providing the leadership and direction required to enable every student to obtain these skills and characteristics prior to graduation. Furthermore, we assert that continuous assessment and modification of curricula and programs is essential, and we believe that these outcomes will aid this process.

Each of the outcomes is presented below in summary form.

Critical Thinkers.

Students will learn to be critical thinkers, demonstrating the ability to define a problem, identify or create solutions, evaluate those solutions, and make an informed decision regarding these alternatives. Emphasis will be placed on encouraging students to ask questions and seek answers related to academic and personal issues.

Good Communicators.

Students will possess strong communication skills, including the ability to both write and speak clearly and effectively, and will be able to use these skills in an appropriate manner when participating in group or person-to-person situations.

Academically Competent.

Students will demonstrate the appropriate level of competency for their abilities and goals in all basic academic disciplines, including math, the sciences, humanities, and technology.

Globally Aware.

Students will demonstrate an awareness of community and an appreciation the diversity of cultures in our society, as well as show respect for the environment and a willingness to participate in government.

Life-long Learners.

Students will develop an attitude of life-long learning, continuing to grow and flourish after they have left the schools. They will respect their bodies and understand the effects of their choices on health and emotional well-being.

Good Workers.

Students will develop a strong sense of respect for themselves and other showing the ability to work positively with others. Important interpersonal skills such as the ability to resolve conflicts, active listening, and the art of negotiation will be encouraged.

Productive.

Students will learn the skills needed to be productive in the workforce, including learning time- and stress-management skills, showing the ability to work both independently and in groups, and be able to adapt to change.

In summary, students will develop a broad set of interests, including creative writing, languages, the humanities, athletics, and a love for reading and will be well equipped to pursue whatever path they choose to take after leaving high school. They will know when and where to seek assistance, how to use language and math productively, demonstrate financial responsibility, and set goals that they can achieve.

Date: Thu Nov 12, 1992 11:33 am PST
Subject: meaning

[From Bill Powers (921112.0900)] Avery Andrews (921112.1238) --

>>But that objection is based on the idea that there is a "real"
>>meaning of "The Iliad" that can be objectively established by
>>experts. All this

>Which there is, sort of, tho `objective' might not be quite the
>right word for. The real Iliad is whatever Homer dictated to
>whoever wrote it down, on current conceptions.

The words "The Iliad" are words. They are not a book, nor are they "a book." "A book" is the name of the thing you can hold in your hands and read. "The Iliad" is a phrase we use to refer to (a) something that we have heard or read that Homer dictated to someone, (b) a particular book with the marks "The Iliad" printed on the cover, and with printed words inside like those that Homer purportedly said, although in most cases not the same words (a translation). So "The Iliad" is the name of a collection of perceptions, some of which are word-like objects, some of which are book-like objects, some of which are memories, and some of which are imagined scenes or people.

When we use a word casually, as in the sentence "My wife is upstairs with my grandson," we don't distinguish between the words and the perceptions that they indicate. The word and its meanings seem inseparable. I speak the words "my wife" and I experience at the same time a perception (in memory/imagination) of a woman I know very well, of memories of other times, of the sound of her voice and the expressions on her face. As I speak the words I simultaneously imagine the things they are about. Of course I also experience other words: "Mary." This is all a sort of silent resonance in memory, with the clearest part being, perhaps, the most recent and relevant to what's going on today, and with little of it being in words.

That sort of resonance in memory occurs with most words that I speak or hear (except for the boilerplate). Seeing the words as separate from their meanings is not normal, I suppose. But then seeing actions as controlled input is not normal either. Normal is what we have learned. We can learn a different way of understanding our experiences of words and meanings.

My understanding of the relation between words and meanings began when I was in high school and ran across Korzybski. I suppose that the maxim "the word is not the object" is taken by a lot of people (judging from the way they fail to apply it) as a sort of moral principle. I took it as Korzybski meant it, I think: as a fact. It is simply an obvious truth that a word is a word and an object is an object. The link between them is invisible. Hearing or seeing the word can make us imagine the object, and seeing the object can make us imagine the word, spoken or printed. Once the difference is noticed, it's hard to remember not noticing it.

While HPCT was under development, I expanded on the basic idea. Not only is the word for an object not the object, but the word for a motion is not the motion and the word for an event is not the event. A description of a relationship is not the perception of the relationship, and so on to saying that the description of a system concept is not the perception of the system concept. Most control processes -- outside language itself -- are involved with direct control of the perceptions, not control of the names or descriptions of perceptions. Words become involved with these control processes at the higher levels, above categories and sequences. I think that linguists will ultimately be the people who tell us just how they are involved. But linguists, too, will have to distinguish the word or description from the "object" it means before this can come about.

Henry Dodson, or Lewis Carroll, seems to have understood this. I can't bring the passage to mind, but perhaps someone can: I think it was the Carpenter who proposed to sing a song called X. But it worked out that X was only what the song was CALLED; its NAME was Y. But the song whose name was Y was actually Z, and the song Z turned out to have a

completely different title. The word is not the object. Nor is the word the meaning of the word.

Bruce Nevin (921112.0915) --

Roger on the line-endings. I have to remember to insert hard carriage returns where they aren't supplied.

RE: outcome-based education.

Looks like a good start toward defining the goals of education in terms of reference perceptions. There seems to be some attempt to break down each one into lower-level reference perceptions, but I think that could be made much more explicit.

" Students will possess strong communication skills, including the ability to both write and speak clearly and effectively, and will be able to use these skills in an appropriate manner when participating in group or person-to-person situations."

In order to be perceived as possessing strong communication skills, what specific sub-goals must the student achieve?. What are the perceptible signs of clearness and effectiveness? What manners are perceived to be appropriate? How do we judge appropriateness?

RE: interpersonal organization.

>there does seem to be some supra-individual collective creation
>that persists over time and is shared across persons.

This has to be a metaphor. Remember that there is no place to stand from which to verify that this is so. Each person has an idea of what is shared and to what extent it is "collective." We go on the assumption that we share a reality. There are limits to our ability to make this assumption believable -- the main limit coming from our knowledge that it IS an assumption.

>The consensus-perceptions are felt to be different from the private
>perceptions, and indeed they are unless one happens to agree with the
>consensus utterly.

How do you know what the consensus is but by perceiving it privately?

Best to all, Bill P.

Date: Thu Nov 12, 1992 3:54 pm PST
Subject: charles dodgson's white knight

THROUGH THE LOOKING GLASS

by LEWIS CARROLL

THE MILENNIUM FULCRUM EDITION (C)1991

[...]

CHAPTER VIII

`It's my own Invention'

[...]

`You are sad,' the Knight said in an anxious tone: `let me sing you a song to comfort you.'

`Is it very long?' Alice asked, for she had heard a good deal of poetry that day.

`It's long,' said the Knight, `but very, VERY beautiful. Everybody that hears me sing it -- either it brings the TEARS into their eyes, or else -- '

`Or else what?' said Alice, for the Knight had made a sudden pause.

`Or else it doesn't, you know. The name of the song is called "HADDOCKS' EYES."'

`Oh, that's the name of the song, is it?' Alice said, trying to feel interested.

`No, you don't understand,' the Knight said, looking a little vexed. `That's what the name is CALLED. The name really IS "THE AGED AGED MAN."'

`Then I ought to have said "That's what the SONG is called"?' Alice corrected herself.

`No, you oughtn't: that's quite another thing! The SONG is called "WAYS AND MEANS": but that's only what it's CALLED, you know!'

`Well, what IS the song, then?' said Alice, who was by this time completely bewildered.

`I was coming to that,' the Knight said. `The song really IS "A-SITTING ON A GATE": and the tune's my own invention.'

Date: Thu Nov 12, 1992 4:08 pm PST
Subject: Revised Intro. Manipulation

[From Dag Forssell (921112)]

For Gary Cziko, Bill Silvert, Mary and others.

The CSGINTRO.DOC posted by Gary yesterday lost the formatting when Gary called up the file from my disk into his Macintosh. Loss of hard returns meant that part of the post got scrambled and lost its clarity.

The best way to get the file to Gary is to post it, so here it is. Gary will post it regularly. Since Gary has asked me to assume responsibility for this document, I will incorporate any sensible suggestions for clarification and updates and then post again.

This posting already incorporates the information on how Bill Silvert created a new subdirectory and a suggestion by Gary to incorporate Usenet address information. With Mary's permission I am incorporating her posting of yesterday and substituting her literature list (I fiddled with it, Mary). I have also merged some of Mary's info in the CSG group section. OK?

Bill Silvert, please substitute this file on the server. I suggest you delete the readme file. It was for you and Gary, not for posterity.

INTRODUCTION TO THE CONTROL SYSTEMS GROUP NETWORK (CSGnet)
AND TO THE CONTROL SYSTEMS GROUP
Prepared by Dag Forssell 921112.

This introduction provides information about:

Why you might want to join the CSG and read CSGnet
Our subject: The control paradigm and Perceptual Control Theory
The evolution of the control paradigm
Demonstrating the Phenomenon of Perceptual Control
The purpose of CSGnet
CSGnet participants
How to ask effective questions
The Control Systems Group
Subscribing to CSGnet
How to obtain text and program files
Literature references

WHY YOU MIGHT READ THE CSGnet

If you have an interest in how people and organisms work.
If you are dissatisfied with the explanations in many of the life sciences and would like a more rigorous approach....
If you insist on thinking things through for yourself....

OUR SUBJECT: THE CONTROL PARADIGM AND PERCEPTUAL CONTROL THEORY

Human control is the primary subject of CSGnet, but all forms of control are game. Here are two introductions, first by Bill and then by Mary Powers:

There have been two paradigms in the behavioral sciences since the 1600's. One was the idea that events impinging on organisms make them behave as they do. The other, which was developed in the 1930's, is PERCEPTUAL CONTROL THEORY (PCT). Perceptual Control Theory explains how organisms control what happens to them. This means all organisms from the amoeba to humankind. It explains why one organism can't control another without physical violence. It explains why people deprived of any major part of their ability to control soon become dysfunctional, lose interest in life, pine away and die. It explains what a goal is, how goals relate to action, how action affects perceptions and how perceptions define the reality in which we live and move and have our being. Perceptual Control Theory is the first scientific theory that can handle all these phenomena within a single, testable concept of how living systems work.

William T. Powers, November 3, 1991

PERCEPTUAL CONTROL THEORY

While the existence of control mechanisms and processes (such as feedback) in living systems is generally recognized, the implications of control organization go far beyond what is generally accepted. We believe that a fundamental characteristic of organisms is their ability to control; that they are, in fact, living control systems. To distinguish this approach from others using some version of control theory but forcing it to fit conventional approaches, we call ours Perceptual Control Theory, or PCT.

PCT requires a major shift in thinking from the traditional approach: that what is controlled is not behavior, but perception. Modelling behavior as a dependent variable, as a response to stimuli, provides no explanation for the phenomenon of achieving consistent ends through varying means, and requires an extensive use of statistics to achieve modest (to the point of meaningless) correlations. Attempts to model behavior as planned and computed output can be demonstrated to require levels of precise calculation that are unobtainable in a physical system, and impossible in a real environment that is changing from one moment to the next. The PCT model views behavior as the means by which a perceived state of affairs is brought to and maintained at a reference state. This approach provides a physically plausible explanation for the consistency of outcomes and the variability of means.

The PCT model has been used to simulate phenomena as diverse as bacterial chemotaxis, tracking a target, and behavior in crowds. In its elaborated form, a hierarchy of perceptual control systems (HPCT), it has lent itself to a computer simulation of tracking, including learning to track, and to new approaches to education, management, and psychotherapy.

Control systems are not new in the life sciences. However, numerous misapprehensions exist, passed down from what was learned about control theory by non-engineers 40 or 50 years ago without further reference to newer developments or correction of initial misunderstandings. References in the literature to the desirability of positive feedback and the assertion that systems with feedback are slower than S-R systems are simply false, and concerns about stability are unfounded.

The primary barrier to the adoption of PCT concepts is the belief - or hope - that control theory can simply be absorbed into the mainstream life sciences without disturbing the status quo. It is very hard to believe that one's training and life work, and that of one's mentors, and their mentors, must be fundamentally revised. Therefore, PCT appeals to those who feel some dissatisfaction with the status quo, or who are attracted to the idea of a generative model with broad application throughout the life sciences (plus AI and robotics). There are very few people working in PCT research. Much of its promise is still simply promise, and it meets resistance from all sides. It is frustrating but also tremendously exciting to be a part of the group who believe that they are participating in the birth of a true science of life.

Mary Powers, November 1992

THE EVOLUTION OF THE CONTROL PARADIGM

The PCT paradigm originates in 1927, when an engineer named Harold Black invented the negative feedback amplifier, which is a control device. This invention led to the development of purposeful machines. Purposeful machines have built-in intent to achieve consistent ends by variable means under changing conditions.

The discovery and formalization of the phenomenon of control is the first alternative to the cause-effect perspective ever proposed in any science.

The first discussion of purposeful machines and people came in 1943 in a paper called: Behavior, Purpose and Teleology by Rosenblueth, Wiener and Bigelow. This paper also argued that purpose belongs in science as a real phenomenon in the present. Purpose does not mean that somehow the future influences the present.

The first specific suggestion on how to use the concept of control to understand people came in 1957 in a paper entitled: A General Feedback Theory of Human Behavior by McFarland, Powers and Clark.

In 1973 William T. (Bill) Powers published a seminal book called "Behavior: the Control of Perception," which still is the major reference for PCT. See literature below.

This book spells out a complete model of how the human brain and nervous system works like a living perceptual control system. Our brain can be viewed as a system that controls its own perceptions. This view suggests explanations for many previously mysterious aspects of how people interact with their world.

Perceptual Control Theory has been accepted by independently thinking psychologists, scientists and other interested people. The result is that an association has been formed (the Control System Group), several books published, this CSGnet set up and that at latest count 16 professors are teaching PCT in American universities today.

DEMONSTRATING THE PHENOMENON OF CONTROL

The phenomenon of control is largely unrecognized in science today. It is not well understood in important aspects even by many control engineers. Yet the phenomenon of control, when it is recognized and understood, provides a powerful enhancement to scientific perspectives.

It is essential to recognize that this phenomenon exists and deserves an explanation before any of the discourse on CSGnet will make sense.

Please download the introductory demonstration demla.exe, which is an interactive program and/or rubberbd.txt, which is a text telling you how to demonstrate the phenomenon to yourself and a friend using only two rubber bands.

THE PURPOSE OF CSGnet:

CSGnet provides a forum for development, applications and testing of PCT.

CSGnet PARTICIPANTS

Many interests and backgrounds are represented here. Psychology, Sociology, Linguistics, Artificial Intelligence, Robotics, Social Work, Neurology, Modeling and Testing. All are represented and discussed. A quality of participants on this net is that most are prepared to question and re-consider what they think they know.

HOW TO ASK EFFECTIVE QUESTIONS

Since PCT puts much conventional, well established wisdom on its head, it is helpful to begin by demonstrating the phenomenon of control to yourself and studying a few references. As you catch on, read this net and follow a thread that interests you for a month or more.

When you ask a question, please consider that in order to give you a good answer, a respondent will need to put your question in context.

Therefore, please introduce yourself with a statement of your professional interests and background. It will help if you spell out which demonstrations, introductory papers and references you have taken the time to digest and what you learned.

People on this net are in various stages of learning and understanding PCT. When you get a reply to your post, please consider that the respondent who found your question of interest and invested time in a reply, may benefit from knowing how you perceived the answer. Did it answer your question? Was it clear? Were you able to understand it?

THE CONTROL SYSTEMS GROUP

The CSG is an organization of people in the behavioral, social, and life sciences who see the potential in PCT for increased understanding in their own fields and for the unification of diverse and fragmented specialties.

Annual dues are \$45 for regular members and \$5 for students (subsidized).

An annual meeting is held in Durango, Colorado, on the campus of Fort Lewis College. In 1993 it will begin in the evening of Wednesday, July 28, and end Sunday morning, Aug, 1. There will be 7 plenary meetings (mornings and evenings), with afternoons, mealtimes, and late night free for further discussion or recreation. Full details will be available on the net or by mail after April 1, 1993.

Net subscribers find it useful to have thematic collections of some of the network discussions, and it enables non-net members to keep up with them. Threads from this net are published on a quarterly basis in a booklet called the Closed Loop. These booklets are distributed to members and are available separately. A complimentary copy of Closed Loop will be sent upon request.

Back issues are available: Volume 1 (4 issues) is \$12. Single issues of Volume 2, beginning with Jan. 1992, are \$6 each.

For membership information and back issues of Closed Loop, write: CSG, c/o Mary Powers, 73 Ridge Place CR 510, Durango, CO 81301-8136.

SUBSCRIBING TO CSGnet

When you subscribe to CSGnet, you get this message, CSGINTRO.DOC. But you may have received it from a friend who printed it, seen it on a demodisk, or seen it on Usenet. To subscribe, send a message as follows: (Internet)

To: LISTSERV@VMD.CSO.UIUC.EDU
Message: Subscribe CSG-L Lastname, Firstname, Affiliation, City, State.
help

(Lastname, Firstname, etc is optional commentary).
("help" requests a list of most commonly used commands).
(The Bitnet address is: LISTSERV@UIUCVMD).
(This server is not sensitive to upper or lower case letters).

CSGnet can also be accessed via Usenet (NetNews) where it is listed as the newsgroup bit.listserv.csg-1.

HOW TO OBTAIN TEXT AND PROGRAM FILES

A number of ASCII documents and binary computer programs are available on a fileserv maintained by Bill Silvert. It is possible to download all these files via e-mail. If you are on internet, it is easiest to obtain binary program files via anonymous FTP. If you are on MCI mail, you have read about how you can transfer binary files with Kermit or Zmodem protocols. (Type help at the MCI mail prompt for directions). But the server cannot send binary files over the internet mail network, so download uue.scr first, then request the binary files uuencoded as ASCII files. The Internet address for the server is BIOME.BIO.NS.CA. CSGnet files are kept in the subdirectory pub/csg.

To get basic information and a current listing of available documents, send a message as follows: (Internet)

To: SERVER@BIOME.BIO.NS.CA.

Message: help
ftp
get csg/Index
end

"help" asks the server to send you commands and explanations.
"ftp" requests the scoop on anonymous FTP for internet.
"get csg/Index" requests the Index for the csg subdirectory.

In your message, pay attention to upper and lower case! DOS is not dos.

As part of the index (of the csg directory), you may be looking at:

programs/msdos:

demla.exe 128437 Bill Power's demonstration of perceptual control
dem2a.exe 123649 Bill Power's modelling of control

documents/forssell:

uud.scr 53406 Create a PC uuencode binary. (Instructions @ end)!

If you want to request demla.exe (uuencoded) to get a "live" demonstration of the phenomenon of control, and the ASCII file uud.scr which allows you to use DOS debug to create uud.exe to decode it, send the following message:

```
uue csg/programs/msdos/demla.exe
get csg/programs/forssell/uud.scr
```

The uuencoded demla.exe will be sent in four parts. Remove headers and use an editor to make it into one file (starting with table and ending with end) before you use uud.exe to restore the file. demla.exe is a self-extracting archive file. Put it in it's own directory before you execute it. You get complete documentation and a running program.

LITERATURE REFERENCES

For a complete list of CSG-related publications, get the file biblio.pct from the fileserver as described above. Here are some selected books and papers on Perceptual Control Theory:

Powers, William T., Behavior: The Control of Perception.
Hawthorne, NY: Aldine DeGruyter, 1973.

The basic text.

*Robertson, Richard J. and Powers, William T., editors.
Introduction to Modern Psychology: the Control Theory View.
CSG Books, 1990. (\$25 postpaid)
College-level text.

*Powers, William T., Living Control Systems I : selected papers.
CSG Books, 1989. (\$16.50 postpaid)
Previously published papers, 1960-1988.

*Powers, William T., Living Control Systems II : selected papers.
CSG Books, 1992. (\$22 postpaid)
Previously unpublished papers, 1959-1990.

*Marken, Richard S, Mind Readings : experimental studies of purpose.
CSG Books, 1992. (\$18 postpaid)
Research papers exploring control.

Marken, Richard S., editor. Purposeful Behavior : the control theory approach. Thousand Oaks, CA: Sage Publications: American Behavioral Scientist, special issue. Vol. 34, Number 1. September/October 1990.
11 articles on control theory.

Runkel, Philip J., Casting Nets and Testing Specimens. NY: Praeger, 1990.

When statistics are appropriate; when models are required.

Hershberger, Wayne, editor, Volitional Action; conation and control. Advances in Psychology 62. NY: North-Holland, 1989.

25 articles (not all PCT)

Ford, Edward E., Freedom From Stress. Scottsdale AZ: Brandt Publishing, 1989.

A self-help book. PCT in a counseling framework.

Gibbons, Hugh, The death of Jeffrey Stapleton: exploring the way lawyers think. Concord, NH: Franklin Pierce Law Center, 1990.

A text for law students using control theory.

McClelland, Kent, Perceptual Control and Sociological Theory. 1992.

Not yet published. Available from the author,
Grinnell University, Grinnell, Iowa.

McPhail, Clark, The Myth of the Madding Crowd. NY: Aldine de Gruyter, 1990.

Introduces control theory as a way to explain group behavior.

McPhail, Clark, Powers, William T., and Tucker, Charles W.,
Simulating Individual and Collective Action In Temporary Gatherings.
Social Science Computer Review, 10:1, Spring 1992.

Durham, NC, Duke University, 1992.

Computer simulation of control systems in groups.

Petrie, Hugh G., Dilemma of Inquiry and Learning. U of Chicago press, 1981.

Introduces PCT to educational theory.

Richardson, George P., Feedback Thought in Social Science and Systems Theory.
Univ. of Pennsylvania Press, 1991.

A review of systems thinking, including PCT.

*These items are available from CSG Books, 460 Black Lick Road,
Gravel Switch, KY, 40328.

--End of document--

Re: Manipulation

While I am paying postage for a fat letter already, let me slip this piece in. I found it amusing.

This is a 2/3 extract from an article by G. Gordon Liddy (of Watergate fame) in a supplement to Forbes magazine called FYI, Nov 23, 1992.

.....
..."Hold it," Bernie commanded, looking as if I had just desecrated a temple.
"You're talking about _fraud_, not a _con_."

"What's the difference?"

"Ever hear the phrase, 'You can't cheat an honest man?' Well, that's as wrong as 'Crime doesn't pay,' and 'Speed kills.' If speed killed, every

Indianapolis driver would be dead. If crime didn't pay, there'd be no crime, and honest men are defrauded every day. The expression should be, 'You can't _con_ an honest man.'

"What's the difference?"

Bernie took on the air of a professor of theology at a seminary. He was teaching the principles of a subject in which he believed deeply. "You can misrepresent yourself to an honest man; promise to perform a service or provide goods for a fair price, then renege--purposely fail to deliver and keep his money--and take off with it. You've _cheated_ him, committed a fraud. But you haven't conned him. To do that you have to appeal to his greed--convince him that _he's_ cheating someone else. _That's_ a con."

Bernie was really into it now. The years dropped off him magically as he relived his youth. "For instance, when I was playing the chauffeur of The Yellow Kid and we were running 'The Eye.'"

"Why," I asked, "was he called 'The Yellow Kid'?"

"Because he always dressed in yellow. It was his favorite color. His trademark: yellow suit, yellow tie, yellow gloves, even yellow shoes."

"Wait a minute," I objected. "He'd be spotted in a minute. How could you run a con looking like that? The marks would think he was a nut."

Bernie looked pained, as if I were a none-too-bright student. "He didn't dress that way when he was _working_. Those were only his leisure clothes. When he worked he dressed the part he was playing--usually a rich businessman. Like when we would run 'The Eye.'"

"Tell me how it worked."

"We'd tour the real rube country in the Midwest in a big Duesenberg or a Rolls limousine, " Bernie grinned. "The Yellow would be sitting in the back. I'm driving. He's wearing clothes worth half the price of the car. I've got on the uniform, with boots and the little peaked cap, see, when we pull up in front of the local bank. I hop out and open the door for The Yellow, and he makes this entrance into the bank and goes to the teller cage to break a couple of \$1,000 bills for change. In the middle of the transaction, he suddenly claps his hand over his eye and half shouts--real agitated--'My eye! My eye!' Then he slips on a pair of dark glasses, still making a lot of noise.

"Now," said Bernie, "the president of the bank comes running out of his office, he sees the Rolls, the clothes, the thousand-dollar bills, and he's terrified this super-rich guy's somehow had his eye put out in his bank. You've got no idea how relieved he is to hear The Yellow reassure him that he lost his eye in an accident years before. What's happened, says The Yellow, is that his _glass_ eye, one made especially for him in Czechoslovakia to perfectly match his good eye and cost a fortune, has popped out and is somewhere on the floor.

"Naturally," Bernie continued, "everybody's on the floor looking, but nobody can find it. Desperate, The Yellow asks, 'Where's the local newspaper office?"

I'll give \$2,000 -no, _\$3,000_- for the return of my eye.'

"Well, of course," said Bernie, "the bank president doesn't want him to do that, he wants to find it. So he asks The Yellow please not to do that; there'll be a thousand people in here looking for it-bad for business. The Yellow says he understands, but that he can't stay. He explains that if they find the eye, 'Call this number, tell the man who answers the phone you have it, and he'll arrive on the next train to claim it and pay the \$3,000 reward.'

"The next day," said Bernie, still professorial but clearly enjoying himself, "a guy walks into the bank to cash a travelers check, goes to the same window The Yellow did, bends over, reaches for something and says, 'Look--I found a marble--only it's not round, it's oblong. It looks more like...a glass eye!_'

"Now the president sees this," said Bernie, "and he says, 'You're right, sir. A man lost it here the other day. He may be willing to give a small reward.' 'Call him,' says the guy who 'found' the eye. 'Whatever I get, I'll split.'

"The president frowns and says, 'One problem. He didn't give me a number to call him. He said he'd call me. Let me check to see if he's called yet.' Then he goes into his office and calls the number. When the guy on the other end answers, the president says, 'I've found the eye. Is that \$3,000 reward offer really good or just conversation?'

Bernie was pacing up and down now, excited. "'Sir,' the guy on the phone says, 'Mr. X always keeps his word.' He pauses and then says back to the president, 'Sir, if I can give you some information worth a lot of money to you, can I get \$500?' Then he tells the president, 'Listen, that eye cost him \$15,000. He'll be glad to give you \$5,000 or even \$6,000.' The president promises the guy on the phone \$500, then goes back out to the guy holding the glass eye and says, 'He hasn't called yet, and I don't know if he ever will. Tell you what; I'll give you \$500 for your end and take my chances.'"

I wish you could have seen the glee in Bernie's face as he continued: "The guy with the eye says no he wants to wait and talk to the owner himself. The president offers \$800. The guy still says no. The president goes up each time the guy refuses to go for the deal, all the way up to \$2,000. Finally the guy with the eye says, 'Look, you offered me \$2,000, I can't wait around. Give me \$2,300.' The president is delighted to give it to him." Bernie chuckled, "And he's still waiting to collect his six grand."

Bernie reverted to his professorial demeanor. "You see? That guy beats himself. If he doesn't get greedy and try to cheat the other guy, he wouldn't be out a dime. Now that's a con."

It sure is.

-END-

Fraud and Con: A distinction to be made as manipulation is codified by Control Theorists!

Best to all, Dag

Date: Thu Nov 12, 1992 4:22 pm PST
Subject: Re: meaning, reality

[Avery Andrews 921113.957] (Bill Powers (921112.0900))

>I think that linguists will ultimately be the people who tell
>us just how they are involved. But linguists, too, will have to
>distinguish the word or description from the "object" it means before
>this can come about.

They already do, as do all philosophers, e.g. Barwise and Perry in their 'situation semantics'. I just don't see what problem you're getting at here. I would say that objective reality should be accepted because it's the best explanation of the experiences that support congruence. E.g. if you tell me that if I follow a certain route I'll find something at the end, & I follow the route and find what you told me I'd find. There are various other possible theories, such as that I am a brain in a vat being fed sense data by a vast computer, but by normal scientific standards, these are obviously hopeless.

Avery.Andrews@anu.edu.au

Date: Thu Nov 12, 1992 4:38 pm PST
Subject: Re: consensual reality, negative evidence, parry/lord

[Avery Andrews (921112.951)] (Bruce Nevin (Thu 921112 07:25:24))

On membership, etc. Yes, I agree, & would only add that there's sometimes more than in-groupiness involved in what concepts are constructed consensually. E.g. if a bunch of people have a suboptimal concept of plutonium they have a good chance of winding up prematurely dead.

Re also you remarks on negative evidence, etc. I agree that I overstated things a bit in saying that my story was the alternative to GB, but it does seem to me as a parent that neither I nor my wife spent much time correcting syntax and subcategorization errors - what stands out is the corrections of morphology, and some usage issues, such as the much/many distinction.

Parry/Lord aren't exactly passee, but people don't buy everything they said anymore. In particular, the formula system, tho obviously present, is probably not as pervasive as they thought it was, and it is also not necessary to believe that the stuff was all made up extempore on the spot - oral traditions vary greatly in the extent to which things are rehearsed or composed beforehand, it is perfectly possible that the bards did quite a lot of composition and rehearsal (in their heads, of course) when they weren't performing. & of course preclassical Greece wasn't necessarily like the Balkans in the 20's.

Avery.Andrews@anu.edu.au

Date: Thu Nov 12, 1992 4:47 pm PST
Subject: Re: o tell me where is meaning bred

(penni sibun 921112.1600) Bruce Nevin (Wed 921111 10:59:45)

No, I am not saying that meaning is everything, I am saying that each

thing "has" specific meanings. Understand that every "thing" is a perception, and each meaning is another perception. I am not saying that a given thing/perception "has property x" i.e. has meaning in some vague, generalized sense of "the property of meaning." I am saying that each given thing-perception has associated with it specific other thing/perceptions, differently for each

as stated, this is either circular or leads to infinite regress. i think perhaps you're trying to say that meaning inheres in the structure of relations bet. perceptions. that's fine, but then i wonder why it's necessary to call it ``meaning'' rather than ``structure'' or ``associations'' or something like that.

1

A word W "means" certain other perceptions for a given person at a given time.

....

Verbal perceptions are distinguished from most nonverbal perceptions in that they participate in a conventionalized structure by which people create and transmit information. There are iconographic systems that are not so articulate, and there are language-like artifacts (notably mathematics and logic) that lack the flexibility and informational capacity of language, which are also distinct from other nonverbal perceptions insofar as they are structured according to agreed conventions.

i think you are, roughly, equating words w/ verbal perceptions. correct me if i'm wrong! although you are obviously attentive to the social context of language, i don't see your proposal as any different from the venerable view of language as encoding/decoding meanings from my head to yours and vice versa (you say ``create and transmit information''; btw, how does info differ from meaning? what is it created out of?). i just don't see how your distinction can fly. i'll take two examples, a verbal one where the ``c&t of info'' is about nil, and a nonverbal in which a lot of info is c&t'd.

i'm the kind of person who regularly panics in a party situation where i'm supposed to mingle. when i'm confronted w/ random strangers, no matter how nice and interesting they probably are, i have nothing to say to them. but, i am strongly constrained to vocalize and, preferably, verbalize. so i (quite incompetently) use all the strategies i know: i mutter typical nonce phrases; repeat the last statement i heard; try to find a productive standard open-ended question; etc.. verbal behaviour, sometimes lots of it, containing almost no information. (arguably lots of info from other behaviour.)

i occasionally find myself in situations more pleasant than parties, such as when i am flirting w/ someone. i'll define flirting as a practice by people who don't know each other well but are in the process of defining (or redefining) their relationship. flirting is mostly or entirely nonverbal, yet an enormous amount of info is c&t'd. flirting does not simply establish, eg, sexual attraction, but communicates all sorts of complex messages such as ``i enjoy x,'' ``don't do y,'' ``this is how that makes me feel.''

i don't care whether you talk in terms of actions or perceptions, i don't think your verbal/nonverbal distinctions have very firm basis in ``reality.''

you are probably familiar with it, but i recommend the appendix in roy harris's the language machine for a mordant discussion of language as encoding/decoding of meaning.

cheers. --penni

Date: Thu Nov 12, 1992 4:58 pm PST
Subject: lousy argument form

[Avery Andrews 191112.1013]

Recently on csgnet I've seen instances of a really lousy argument, to wit, X doesn't understand PCT because they're controlling for believing something else. Aside from the fact that this kind of argument can also be used to explain why people think that PCT is true (you guys believe in this stuff because you want to think you have some control over what you do, rather than the truth, which is that you are the idle playthings of Zeus & Co. on Mt. Olympus), but also it should be evident that anyone who bothers to get CSGNet is not committed to believing the standard stories being told in their discipline. After all, even deleting the traffic without reading it is a substantial investment of effort.

Avery.Andrews@anu.edu.au

Date: Thu Nov 12, 1992 5:11 pm PST
Subject: Re: o tell me where is meaning bred

The distances are such that they can hop it with no serious problems, I think.

Surprised not to have heard much for forsell & folson -- maybe my little blasts are having some effect. Almost reach the flame level today!

Date: Thu Nov 12, 1992 6:14 pm PST
Subject: More on Hamming

From Bill Cunningham (921112.2030)

Re John Gabriel (921111.0945), his posts on Hamming distance, Bruce's comment.

First, John's intro failed to state the official method for finding the needles in haystacks: Process the haystacks through the cattle herd, then study the resultant product in great detail. Since that approach is fun, add more haystacks and cattle--but never, never, never use a magnet, look for sick cows, or even the guy hiding the needles in the first place. You may conclude we were not successful. That provides our current motivation.

Comments on Hamming distance are quite in order. Actually, had only considered two protagonists--one of whom serves as observer. First, both players have to control for communicating--at least for now. A Hamming distance can only be defined for the mutually held proposition set. No real progress is possible unless both parties agree to make this the union of their respective sets, with the disjoint members weighted at zero by one party or the other. How do they get this far? First layer of layered protocols should to establish the larger set. Subject to error? Sure. Recursive? You bet. But Martin Taylor's point that communication is the control of belief applies. The recursion should continue until both are satisfied the larger set is mutually held. Presumably both parties will continue to test for zero perceived error on this point.

Having gone through all this, a Hamming distance can be defined since a difference in individual proposition weights is only difference between the proposition sets held by the two protagonists.

All of this merely placing a formalism on LPT.

It's not hard to envision (or explain) various intentional and unintentional dysfunctional approaches to communication (eg deliberate ambiguity, deception, refusal to listen).

John's comment about a row with his boss was stimulated by my observation that many disagreements were fueled by failure to generate the union set as the first protocol layer. The LPT recursion can deal only with the mutual set of propositions. Unfortunately, if each party associates these with other (privately held) propositions, each recursion can only produce a major error signal to each party. The result: "YOU'RE NOT LISTENING TO ME!!" In other words, hidden propositions are downright disruptive, whether intentionally hidden or not.

Regards, Bill C.

Date: Fri Nov 13, 1992 3:49 am PST
Subject: meaning

[Avery Andrews 921113.1047] (Bill Powers (921112.0900))

>"The Iliad" is a phrase we use to refer to (a) something that we
>have heard or read that Homer dictated to someone, (b) a particular book
>with the marks "The Iliad" printed on the cover, and with printed words
>inside like those that Homer purportedly said, although in most cases
>not the same words (a translation). So "The Iliad" is the name of a
>collection of perceptions, some of which are word-like objects, some of
>which are book-like objects, some of which are memories, and some of
>which are imagined scenes or people.

I'm unsure about how all this really gets grounded out in perceptions, considering that the book has to contain what Homer putatively said, probably translated into some other language etc., & that people can use the word successfully without having much of a notion of what is in the book (as when the lit. teacher assigns it as reading to a group of students have never heard of it before).

The consensual element seems to be being left out. Here's an example that might help to fill it in. Aeschylus wrote a play, Seven Against Thebes, which I have not read in any language, & so don't have much of a concept of. Suppose I go to the book store, & buy a book which says on the cover that it is Seven Against Thebes by Aeschylus.

But when I start reading, something seems odd (maybe all the action seems to about a war of independence being fought on the Moon). So I take my book down to the library & compare it with the appropriately titled books there, & find it quite different. I conclude that I've been the deliberate butt of a wierd joke, or an accidental victim of a bookbinder's error.

So, I speculate, part of a normal person's concept of literary works is that there are these places that have `standard copies' of works by given authors. They include libraries, and other collections maintained by supposed experts. How do we know that

they are experts? Mostly, I guess because they say the more or less the same thing, for a wide range of cases, and, partly, because they give a general impression of knowing what they are talking out.

So I'll contend that to get a perceptually based concept of the Iliad we need to invoke (a) some kind of social perception of expertise (b) perceptions of the apparent experts agreeing on whether a given book is a copy of the Iliad. And even people who would trust themselves to recognize copies of the book base their trust on these considerations, except that they encountered the experts first (or just happily assume that their bookseller did gull them once upon a time).

I get the sense that there's some kind of 'second order' perception going on, tho definitely not in the standard PCT sense.

Avery.Andrews@anu.edu.au

=====

Date: Fri Nov 13, 1992 5:39 am PST

Subject: Johari & proposition sets

From Bill Cunningham (921113.0830)

It strikes me that the mutually held proposition set is analogous to the "arena" portion of the Johari window, with the same problems created by proposition sets held unilaterally. I'm not sure how far to carry the analogy, but it might be a useful tool for explaining the mathematical approach to a nonmath oriented audience. If so, the analogy serves as an example of creating mutually held propositions as the first protocol layer for communication.

Bill C.

=====

=====

Date: Fri Nov 13, 1992 7:23 am PST

Subject: Avery: consensus

[From: Bruce Nevin (Fri 921113 07:53:30)] Avery Andrews (921112.951)

Re membership & consensual reality, yes of course, we test agreements against "private" (nonconsensual) perceptions. But not all the time or rigorously or exhaustively, that would be too exhausting and would cripple higher purposes. No, we live to an astonishing extent in a universe of imagined perceptions and "the world" lets us get away with it. I think consensual reality enables this: the agreements that are reached and inherited (learned) as part of culture are such as don't get us into trouble. That is their evolutionary value.

>does seem to me as a parent that neither I nor my wife spent much time

>correcting syntax and subcategorization errors - what stands out is the
>corrections of morphology, and some usage issues, such as the much/many
>distinction.

Most correction is done by the child without adults' awareness, comparing what she said with what adults said back, etc., where what is said back etc. (even if providing a corrected image) was not said with the conscious intention of correcting.

I understood from Parry & Lord when I read their book years ago that the epics certainly must have been composed and rehearsed beforehand, not extemporized on the spot, and that the repeated formulaic passages helped them to remember it all, and gave them a breather of sorts to get set for the next section of more demanding material. My observation is that just this sort of thing happens a lot in conversations--Penni just gave a sort of limiting case where no sort of conversation is foreseen (in imagination) that is not too demanding, and nothing is produced but breather-formulae. It certainly happens a lot in politicians' Q&A responses.

(Avery Andrews 921113.1047) --

The perceptions of agreements that we each build up internally to constitute consensual reality for ourselves do seem to be out there in the world, just as other perceptions seem to be out there in the world. (Bill: as before, the operative word here is "seem".)

Bruce bn@bbn.com

=====
Date: Fri Nov 13, 1992 7:23 am PST
Subject: Penni: information

[From: Bruce Nevin (Fri 921113 09:52:14)] (penni sibun 921112.1600) --

Thanks for posting the quote from Alice. It saved me typing it! (You must have an on-line copy. Is that available someplace?)

> [From: Bruce Nevin (Wed 921111 10:59:45)]

> No, I am not saying that meaning is everything, I am saying that each
> thing "has" specific meanings. Understand that every "thing" is a
> perception, and each meaning is another perception. I am not saying
> that a given thing/perception "has property x" i.e. has meaning in some
> vague, generalized sense of "the property of meaning." I am saying that
> each given thing-perception has associated with it specific other
> thing/perceptions, differently for each

> as stated, this is either circular or leads to infinite regress.

It is not circular in the usual sense invoked in discussions of formal semantics, because we have stepped out of the realm of words and symbols. It is circular only to the extent that the associativity of perceptions one with another turns out to be circular, and that simply is as it is. Similarly, there is no infinite regress unless associative memory is a recursive process with self reference, and that seems prima facie doubtful. Even if such recursion were found, possibly there are inhibiting processes that come into play when some associatively evoked perception turns out to be a controlled perception, and possibly failure of that underlies some kinds of mental illness, but that's just off the cuff speculation.

> i think perhaps you're trying to say that meaning inheres in the
> structure of relations bet. perceptions. that's fine, but then i
> wonder why it's necessary to call it ``meaning'' rather than
> ``structure'' or ``associations'' or something like that.

No. I am saying that the ensemble of associated perceptions IS what we call meaning. To reify the word "meaning" is a mistake. To say that perceptions are meaningful to us suggests (because of the way English works) that there is some substance or essence called "meaning" that "inheres in" the perceptions. But there are only the perceptions themselves. To say that a perception is significant for us is only to say that we are controlling that perception--that is, that we maintain in memory a reference (preferred amount) for that perception, that a comparator function within us produces an error signal when the perception and the reference differ, and that we do whatever it takes to reduce that error (assuming that conflict does not conceal it, and assuming that higher level control does not override it, resulting in willing suspension of control despite error).

A structure of relations between perceptions creates information. In particular, the structure of relations between elements of language (language-perceptions) constitutes linguistic information in utterances.

For linguistic information, I refer you to:

Z. Harris. 1990. Language and Information. New York: Columbia University Press.

----- . 1992. A Theory of Language and Information. Oxford: Clarendon Press.

-----, M. Gottfried, T. Ryckman, et al. 1989. The Form of Information in Science: analysis of an immunology sublanguage. Boston Studies in the Philosophy of Science. Dordrecht: D. Reidel.

We also talked about linguistic information here a year or more ago. One could dig some stuff out of CSG archives given time and motivation.

>i think you are, roughly, equating words w/ verbal perceptions.

There is nothing but perceptions available to us. The distinction I make between verbal and nonverbal perceptions has to do with convention, arbitrariness, and participation in constituting linguistic information.

> i don't see your proposal as any different
>from the venerable view of language as encoding/decoding meanings from
>my head to yours and vice versa

In a code, there is an external table of equivalences between elements of the code and elements of language. Linguistic information is in a coded message, modulo substitutions through the table, precisely to the extent that it is in the decoded message. (Perhaps there is some legitimate use of "code" that does not rest upon language, but I have yet to encounter it. I believe that any other use of the concept of encoding is in fact only metaphor.)

There is no external table of equivalences for language. Nor is there anything with which the elements of language can be interconverted by such a table, were it to exist. The closest analog is the association of perceptions with language-perceptions within a person. This is not stable and is many-many rather than one-one. Salience in the associated perceptions (meanings) depends upon control of those perceptions on the occasion, not upon something like a grammar in which things are looked up and correlated. I should think this view would be very congenial with yours, in at least some respects.

To be sure, there are constraints. Language is a matter of learned convention (language-particular) within the general characteristics of perceptual control (language universals). The discussion with Avery about agreements, membership, consensual reality, and so on bears directly on these constraints. That discussion suggests how there can be stabilities without tables for encoding.

Roy Harris has some useful things to say about de Saussure. He seems deeply vested in an "encoding" conception of language. He rails a lot at other linguists. His review of a book by Zellig Harris a year or so ago was exceptionally uncomprehending and self serving, merely a forum for riding his hobby yet again.

In the two social situations that you described, you should be aware that your perception of what you are doing is quite likely different from the perceptions that other participants have. In the first situation in particular you are deliberately and probably effectively dissembling. In the second situation, you and your flirtation partner are very much involved with imagined perceptions. Is it necessary to say that there are many misunderstandings in this process? That it is only a beginning, an assertion of availability and interest, and that it goes nowhere without on subsequent occasions spending some quality time (Ed's definition of the term) together? We call this evocation of perceptions and negotiation of agreements information, of course, but that is an informal, metaphoric use of the word, not the technical sense of linguistic information that I am using.

Bruce bn@bbn.com

Date: Fri Nov 13, 1992 7:29 am PST
Subject: Bill Cunningham: hidden agendas

[From: Bruce Nevin (Fri 921113 09:53:34)] Bill Cunningham (921112.2030)

>Having gone through all this, a Hamming distance can be defined since a
>difference in individual proposition weights is only difference between the
>proposition sets held by the two protagonists.

A Hamming distance can be defined internally for each participant wrt that participant's perceptions of their agreements. (An observer is also a participant.)

> many disagreements were fueled by failure to generate the union
> set as the first protocol layer. The LPT recursion can deal only with
> the mutual set of propositions. Unfortunately, if each party associates
> these with other (privately held) propositions, each recursion can only
> produce a major error signal to each party. The result: "YOU'RE NOT
> LISTENING TO ME!!" In other words, hidden propositions are downright
> disruptive, whether intentionally hidden or not.

There are ALWAYS hidden propositions. Rather, there are always perceptions associated with a given agreement in one person that are not associated with that agreement in the other. An agreement is itself not an agreement about propositions, it is a (perception of) agreement about perceptions. (An agreement that was only about a proposition would be an agreement like "yes, we employed that string of words.") The difficulty of making propositions "binding" is what keeps contracts specialists, lawyers, and the like busily employed. The effect is based upon the fact that some propositions are performatives (they are constitutive). "I pronounce you man and wife." "Let there be peace between our nations from this time forth." But that they should be performatives (when uttered and perhaps signed by some duly constituted representative, ruler, etc.) is itself a function of agreement, as is indeed the "power vested in" the one uttering the performative. Child's play, really.

Bruce bn@bbn.com

Date: Fri Nov 13, 1992 8:57 am PST
Subject: Bill: reality of consensus of reality of THIS

[From: Bruce Nevin (Fri 921113 10:39:51)] Bill Powers (921112.0900)

> RE: interpersonal organization.

> >there does seem to be some supra-individual collective creation
> >that persists over time and is shared across persons.

> This has to be a metaphor. Remember that there is no place to stand from
> which to verify that this is so. Each person has an idea of what is
> shared and to what extent it is "collective." We go on the assumption
> that we share a reality. There are limits to our ability to make this
> assumption believable -- the main limit coming from our knowledge that
> it IS an assumption.

This was my referent when I said (this morning's post) that "seem" is the operative word here. "Seem to us" would have been clearer.

> >The consensus-perceptions are felt to be different from the private
> >perceptions, and indeed they are unless one happens to agree with the
> >consensus utterly.

> How do you know what the consensus is but by perceiving it privately?

Yes, then "private perceptions" is a misleading term. Substitute "perceptions not subjected to test for consensus" or something similar.

In this we (you and I right now) are of course negotiating agreements and consensus. Has such an agreement any status apart from an individual reader's private perceptions? I suggest that there are prior agreements that it does, such as conversational implicatures and the conventions by which an utterance may constitute an instance of a speech act, and I believe that you and I both participate authentically in such agreements.

Does agreeing that it is so make it so? Just how much magical force do constitutive acts have? Or, to put it the other way round, on what grounds can one postulate that physical reality is a Ding an sich "out there" in Boss Reality, but that social reality is not?

Presumably because we perceive the latter as more malleable, variable, apparently constituted for the sake of cooperative action, and we perceive the former as universal and changeless, at least at the level of physical and chemical regularities or "laws". But in historical perspective science seems just as much a constitutive cultural phenomenon, and several forms of quantum physics declare the constitutive role of the observer. A weak argument, it seems to me.

No, the question of veridical correspondence between perception and perceived seems to me equally moot for both cases. So all we have are perceptions, among which are agreements that we shall act as if the world is as we perceive it in both its physical and its social aspects.

To decline these fundamental agreements is too costly, in terms of personal ability to control in a social world. A third option, one recommended by people like Meister Eckhard and Chuang Tzu, is to participate as though agreeing, while knowing better. The chores are still there to be done, but this inward reversal does inform one's way of chopping wood and hauling water in remarkable ways.

Speaking of which, I have some chores to do, and hopefully my account on the development system for the T/10 project is set up so I can grab those source files now. I will act on the assumption that those objects and associated commitments and agreements are really, really there.

Bruce bn@bbn.com

Date: Fri Nov 13, 1992 7:29 am PST
Subject: Bill Cunningham: hidden agendas

[From: Bruce Nevin (Fri 921113 09:53:34)] Bill Cunningham (921112.2030)

>Having gone through all this, a Hamming distance can be defined since a
>difference in individual proposition weights is only difference between the
>proposition sets held by the two protagonists.

A Hamming distance can be defined internally for each participant wrt that participant's perceptions of their agreements. (An observer is also a participant.)

> many disagreements were fueled by failure to generate the union
> set as the first protocol layer. The LPT recursion can deal only with
> the mutual set of propositions. Unfortunately, if each party associates
> these with other (privately held) propositions, each recursion can only
> produce a major error signal to each party. The result: "YOU'RE NOT
> LISTENING TO ME!!" In other words, hidden propositions are downright
> disruptive, whether intentionally hidden or not.

There are ALWAYS hidden propositions. Rather, there are always perceptions associated with a given agreement in one person that are not associated with that agreement in the other. An agreement is itself not an agreement about propositions, it is a (perception of) agreement about perceptions. (An agreement that was only about a proposition would be an agreement like "yes, we employed that string of words.") The difficulty of making propositions "binding" is what keeps contracts specialists, lawyers, and the like busily employed. The effect is based upon the fact that some propositions are performatives (they are constitutive). "I pronounce you man and wife." "Let there be peace between our nations from this time forth." But that they should be performatives (when uttered and

perhaps signed by some duly constituted representative, ruler, etc.) is itself a function of agreement, as is indeed the "power vested in" the one uttering the performative. Child's play, really.

Bruce bn@bbn.com

Date: Fri Nov 13, 1992 8:57 am PST
Subject: Bill: reality of consensus of reality of THIS

[From: Bruce Nevin (Fri 921113 10:39:51)] Bill Powers (921112.0900)

> RE: interpersonal organization.

> >there does seem to be some supra-individual collective creation
> >that persists over time and is shared across persons.

> This has to be a metaphor. Remember that there is no place to stand from
> which to verify that this is so. Each person has an idea of what is
> shared and to what extent it is "collective." We go on the assumption
> that we share a reality. There are limits to our ability to make this
> assumption believable -- the main limit coming from our knowledge that
> it IS an assumption.

This was my referent when I said (this morning's post) that "seem" is the operative word here. "Seem to us" would have been clearer.

> >The consensus-perceptions are felt to be different from the private
> >perceptions, and indeed they are unless one happens to agree with the
> >consensus utterly.

> How do you know what the consensus is but by perceiving it privately?

Yes, then "private perceptions" is a misleading term. Substitute "perceptions not subjected to test for consensus" or something similar.

In this we (you and I right now) are of course negotiating agreements and consensus. Has such an agreement any status apart from an individual reader's private perceptions? I suggest that there are prior agreements that it does, such as conversational implicatures and the conventions by which an utterance may constitute an instance of a speech act, and I believe that you and I both participate authentically in such agreements.

Does agreeing that it is so make it so? Just how much magical force do constitutive acts have? Or, to put it the other way round, on what grounds can one postulate that physical reality is a Ding an sich "out there" in Boss Reality, but that social reality is not?

Presumably because we perceive the latter as more malleable, variable, apparently constituted for the sake of cooperative action, and we perceive the former as universal and changeless, at least at the level of physical and chemical regularities or "laws". But in historical perspective science seems just as much a constitutive cultural phenomenon, and several forms of quantum physics declare the constitutive role of the observer. A weak argument, it seems to me.

No, the question of veridical correspondence between perception and perceived seems to me equally moot for both cases. So all we have are perceptions, among which are agreements

that we shall act as if the world is as we perceive it in both its physical and its social aspects.

To decline these fundamental agreements is too costly, in terms of personal ability to control in a social world. A third option, one recommended by people like Meister Eckhard and Chuang Tzu, is to participate as though agreeing, while knowing better. The chores are still there to be done, but this inward reversal does inform one's way of chopping wood and hauling water in remarkable ways.

Speaking of which, I have some chores to do, and hopefully my account on the development system for the T/10 project is set up so I can grab those source files now. I will act on the assumption that those objects and associated commitments and agreements are really, really there.

Bruce bn@bbn.com

Date: Fri Nov 13, 1992 9:19 am PST
Subject: o tell me where is meaning bred

[From Rick Marken (921112.0830)]

Just got back to a s**tload of mail. Before I try to comment on some of this stuff, a quick BRAVO to Bruce Nevin for his (Wed 921111 10:59:45) post on meaning. Bruce ends with:

>I am sure that sometime in the future I will be able to put all this
>more clearly and succinctly.

I think you made a hell of a great first stab at it in this post.

Best regards Rick

Date: Fri Nov 13, 1992 10:34 am PST
Subject: MESSAGE FROM MARY

[From Rick Marken (921113.0900)] Mary Powers(921111)

> I ran across a nice quote about the Test. Almost 3 centuries old.

Search then the ruling passion; there alone
The wild are constant, and the cunning known;
The fool consistent, and the false sincere;
Priests, princes, women, no dissemblers here.
This clue once found unravels all the rest.

Alexander Pope, 1688-1744

WOW!!!! Couldn't have said it better myself; "no dissemblers here".
Incredible. Nice find. Mary!!

Best regards Rick

Date: Fri Nov 13, 1992 11:40 am PST
Subject: propositions

[From: Bruce Nevin (Fri 921113 12:34:38)]

(Martin Taylor 921113 11:15) --

> three useful levels of feedback about
>a proposition (possibly a very complex one):
>
>1) to repeat the words back assured the original talker (O) that the
> recipient (R) had received the correct words.

>2) To paraphrase the proposition assured O that R had understood the
> interrelations (syntactic, semantic, pragmatic) of the words, and

Thanks for this about Pask, whom I too have always found needlessly obscure. Henry Hiz also wrote about paraphrasing to demonstrate understanding.

>3) To describe the effect of the proposition in another way (O said "circle"
> and R says "you mean the locus of points equidistant on a plane from a
> given point") means that R understood the proposition.

To "describe the effect of the proposition" is to describe the perceptions with which one associates the proposition, no?

>Pask argued that third-level feedback gave an iron-clad guarantee that
>R had received what O intended to send. Such an agreement seems to me
>to be more than a simple agreement about a simple perception, and far
>more than 'an agreement like "yes, we employed that string of words."'

But (3) is not an agreement about a proposition, unless you mean "proposition" in some prelinguistic sense, e.g. a supposed mentalese. Perhaps I misunderstood what Bill Cunningham meant by the word. If we are interposing some third sort of entity between language perceptions and nonverbal perceptions, please elaborate.

Bruce bn@bbn.com

Date: Fri Nov 13, 1992 12:05 pm PST
Subject: o tell me where is BRAVO bred?

Thanks, Rick! My, that did feel good! Much appreciated.

Bruce bn@bbn.com

Date: Fri Nov 13, 1992 12:24 pm PST
Subject: LEVELS

FROM I.N.KURTZER 921113.1200

to powers,bourbon,hershberger, and other interested parties

i read n article recently where the experimenters "raised kittens in a enviroment illuminated by a strobe light, which flashed eight times a second for ten microseconds each..." the result was that the kittens' neurons that respond to shapes did, but the neurons that respond to motion did not! well, i interpreted this as a verification of a configuration level and a transition level. and since the cat coould detect configurations but could not detect transitions would that put transitions higher in the hiearchy than configurations? neat experiment, needs replication!!

Cynader, M., & Chernenko, G. (1976).
Abolition of direction selectivity
in the visual cortex of the cat.
SCIENCE, 193, 503-505. (8)

in a hurry, bye i.n.kurtzer

oh yeah, what with all these latin quotes? a new rage sweeping the nation or just bored academians?

Date: Fri Nov 13, 1992 1:37 pm PST
Subject: lousy argument form

[From Rick Marken (921113.1000)] Avery Andrews (191112.1013) --

>Recently on csgnet I've seen instances of a really lousy argument, to
>whit, X doesn't understand PCT because they're controlling for
>believing something else.

If the argument is "you can't see the TRUTH of PCT because, according to PCT, you are controlling for not believing PCT" then that IS a VERY lousy argument. I didn't notice anyone making it, though. What I have heard occasionally is the suggestion that maybe some people are having trouble understanding the PCT model because they are controlling for perceiving principles, programs or system concepts that are in- consistent with PCT. This is just offered as a hypothesis to explain phenomena like the following: you explain PCT to a person, showing that control systems control their own perceptions. Then you hear the person say "yes, I understand" and then go off and describe control in SR terms. This is not a rare phenomenon -- why does it happen? My explanation is in terms of PCT, obviously; in order to maintain control of a current system concept, say, the PCT model is perceived in a way that is consistent with the reference for the system concept. This is not a matter of PCT justifying itself; the PCT model (of the person and in general) may be wrong -- that must be TESTED. But I think PCT is a legitimate model of why people have difficulty understanding or accepting ideas that conflict with their own. But it is just a model -- which can make falsifiable predictions of phenomena.

> Aside from the fact that this kind of
>argument can also be used to explain why people think that PCT is true

This is the obvious problem if PCT were being used as an argument to justify itself. But it is something that SHOULD be done if PCT is being offered as a model of human nature; it should apply to the proposers of the model AND to the "proposees". The PCT modeller does not consider him/herself exempt from his/her own model. So if we propose that a person has difficulty understanding PCT because they are defending a system concept, then we (PCT modellers) expect the same to be true of ourselves; some of us (like Bill P.) are better at remembering this than others (like myself). I have no doubt that I may not be understanding the point of view of others as I try to defend my perception of a PCT/

modelling based system concept. But this is why modelling is so important; a lot of these misunderstandings occur because the we are talking about our cherished imaginings. Models (that map into perceptual phenomena) let us check whether these imaginings, cherished or not, can explain our experience -- ie. our perceptions. If there is a model that is not organized as a PCT model that can control then I will believe it when I see it; no problem.

My mind is already boggled by the philosophical implications of all this stuff. The main point, though, is this: if you EVER catch anyone saying "PCT is right and if you don't see that it is you are just controlling a perception that prevents you from seeing this" then turn them in -- this is not science and it's certainly not PCT.

Best regards Rick

Date: Fri Nov 13, 1992 2:31 pm PST

Subject: Meaning etc.

[From Bill Powers (921113.1000)] Penni Sibun (921112) --

Your post of the Knight's Song is doubly welcome. First because you took the trouble to do it for us; second because you made no reference at all to the state of my neurons.

Thanks.

Avery Andrews (921113.0957) --

>I just don't see what problem you're getting at here.
>I would say that objective reality should be accepted because it's
>the best explanation of the experiences that support congruence.

My point is much simpler than the way it's being taken. But it involves two levels of analysis: the relationship between words and (other) perceptions, and the relationship between perceptions and the world. The first one takes place entirely in the brain; the second between the brain and whatever it is that lies outside it.

> ... if you tell me that if I follow a certain route I'll find something
>at the end, & I follow the route and find what you told me I'd find.

I agree that this is how we establish a private sense of a consensual reality. But your example collapses the two levels of analysis into one. First you tell me about a route: this means that you deliver words to me, which I perceive and then interpret into meanings about signposts, terrain, distance, and some image of the something at the end (so I can know when I have found it). I then use these words and the meanings they evoke in me as the basis for constructing reference perceptions about a landscape and a path through it, which I proceed to create in present time by walking. As I walk I perceive many things, only a few of which you mentioned. These are my perceptions, too: I have faith that they correspond to something, both for me and for you, but I can't prove that. When I have moved the landscape and landmarks backward past me far enough, I come to a perception (or a perception comes to me) that matches my interpretation of your description of the thing at the end. The initial match is image-to-image. Then I check to see whether I would describe my perception in the same terms you gave me. Usually there are differences: I might say "Oh, it's a lot smaller than I thought" or squarer or more expensive. But I can see that the discrepancy is due to the meanings I imagined that were not actually in your words; you said "Mercedes Benz" but you neglected to say it was a TOY car. I imagined incorrectly in the dimension of size.

Having revised my imagined perceptions to fit both your words and my real-time perceptions, I now feel that I know what you meant. I don't know whether I've revised them in the right way -- maybe I didn't walk far enough and the toy car wasn't what you meant. But by checking with you I can become much surer that I know am giving the meanings to your words that you had in your head when you said them.

What is still untested, and probably untestable, is whether my perception of the toy car (or of you) corresponds to some form in the world outside my brain, or is a transformation of external variables into something uniquely human. Behind even this question is an article of faith: that there is a stable objective universe out there. We can accept that there is such a stable objective universe without being able to state its exact relationship to the perceptions we experience. We can agree on a description of what each of us perceives to be that common world, but this does not tell us that our perceptions do not BOTH include a transformation of which we are unaware. There may be a "protocol layer" below the lowest layer we recognize.

Even at the level of perceptual connections with words, we can end up with different perceptions and believe that we agree, when in fact our perceptions differ quite a bit. Through continued communication, this process will converge to a state in which neither of us perceives any discrepancies any more. But that is only a state of zero apparent error, and it does not assure that we then have the same perceptions. It merely says that we are satisfied that we do.

This relates to what John Gabriel(921111.0945) and Bill Cunningham (921112.2030) have been talking about. Whether we are speaking of a universe of propositions that is bounded, or the universe of all possible propositions, the perception of the way another person understood our words remains our own perception. The third party who hears both sides is in an even worse position, because now interpretations of TWO verbal inputs are involved. It isn't necessarily true that if C believes that A and B agree, A and B would also say that they agree. The reason is that all three parties have different experiences from which they draw meanings. We now have three pairs of interactions to worry about. And even if all three pairs converge to a mutually-satisfactory conclusion, no one of them knows how the meanings thereby derived relate to the world outside them.

The only way to reach a reliable convergence is to forget about meanings and just look at the logical forms of the propositions, within a simple prescribed set of rules. Maybe Hamming distance can help in this situation. But my hunch is that it won't. I don't think it can help except inside a single person who is comparing two sets of propositions in terms of one private set of meanings. That's what "agreement" boils down to in the end, isn't it?

Bruce provided an example in his list of outcomes for outcome-based education. One of the words used was "appropriate." I think it's clear that three people who agree that a given performance is "appropriate" could each be using a completely different perception of appropriateness, yet seem to agree completely.

We have quite a few different approaches to linguistics on this net. I've been wondering: is there anything about language on which ALL you linguists can agree? It seems to me that as in physics, there has to be a level of description that is so simple and low-level that we can at least agree on what phenomenon is occurring at that level. What are the absolutely non-controversial aspects of language that different people explain differently?

Avery again --

>I'm unsure about how all this really gets grounded out in perceptions,
>considering that the book has to contain what Homer putatively said,
>probably translated into some other language etc., & that people can
>use the word successfully without having much of a notion of what is in
>the book (as when the lit. teacher assigns it as reading to a group of
>students have never heard of it before).

We're talking how brains work, not about whether books "really" contain what Homer said. The way brains are supposed to work in PCT is that they construct inside themselves perceptions that the brain itself interprets as being a real outside world (justifications aside). So all that the brain knows of the world is in the form of perceptions. Brains can argue with each other, by passing words back and forth, about what is really real and truly true about the world. But they're talking about perceptions all the time.

If the point is to establish the authenticity of a copy of the Iliad, then the argument isn't about human brains any more. It's brains talking about what concerns them about the apparent world, which is accepted at face value. PCT gives us a way to step back from such arguments and ask how it is possible that they occur in the first place, and what is going on inside the respective brains as an argument takes place. I don't know whether you'd say this is going "up" a level, but it's certain going to a different point of view in which all discussions about the properties of the external world are only evidence about how the brain works.

Bruce Nevin says (921113.0753)

>we live to an astonishing extent in a universe of imagined
>perceptions and "the world" lets us get away with it.

To which I add only that these imagined perceptions are the main reason that we think we agree with each other. True agreement is incredibly hard to reach, or verify. Perhaps it's impossible.

The latest Closed Loop has arrived. As usual it conceals an incredible job of organizing and editing by Greg Williams -- not to mention the elegant typesetting. Long coherent discourses are generated out of scattered pieces, and you can't see the joints. My own thoughts come back to me in a deceptively well-organized form, for which I am grateful.

Ed Ford --

Along with the Closed Loop, we received your cards on Teaching Responsible Thinking. I think you need to take a closer look at what you're recommending, or at least the way you put it.

The problem, which I'm not the first to notice, is that this approach presumes that the people to whom it is applied are already under control by the use of socially-sanctioned force. It assumes that one person is in a position to set rules and standards for another. When you ask "What are the rules," you're asking not for the rules that the other sets for himself or herself, but the rules that someone else is in a position to enforce by any means necessary. So you're not really offering any choices at the higher levels of organization. I think that it would be difficult to teach a mature and understanding attitude toward rules when there is no choice about them.

As you explain it, this method is best suited for institutional work where the inmates have freedom only within very narrow boundaries, and are not permitted to do anything they want, or to want everything that makes sense to them. "The rules" constitute a rigid framework that constrains everything else and they come from outside the person, not from inside. I don't think that simply setting up the rules as absolutes or givens will teach responsible thinking. It teaches responsibility TO but not responsibility FOR.

Perhaps you would find some new approach if you thought more in terms of teaching responsibility to people who are free to leave anytime they like. They're still subject to social rules and social coercion, but not anywhere nearly to the extent that they would be in an institution. When you ask free people "what's the rule?" you don't have the power to come down on them when they say they don't like the rule. Instead, you have to talk about it, and ask them what rules they would prefer, how they would put it all together. I think that the same approach you use in helping people prioritize their goals would work with rules: what rules could you accept? What rules make sense, for you and everyone? How would YOU set it up? How do you plan to handle rules you don't agree with so your life can still work out in your favor? What rules would you be willing to try out?

Maybe you should think about approaching people in institutions as you would people outside them. Teach them how to use PCT in the real world they'll be going back to, not the artificial and cramped world they're in at the moment.

Best to all, Bill P.

Date: Fri Nov 13, 1992 4:00 pm PST
Subject: lousy argument form

[Avery Andrews 921114.7859] (Rick Marken (921113.1000))

Fine. More dramatic overstatement on my part, I guess.

But of course the idea that people tend to squeeze new ideas into the procrustean bed of old ones is not original with PCT. My guess as to what's hard about it is that although it seems childishly simple as a doctrine, the actual implications tend to be surprising (similar in a way to the legendary executioner who cut people's heads off so deftly that they didn't notice until they tried nodding).

Avery.Andrews@anu.edu.au

Date: Fri Nov 13, 1992 4:09 pm PST
Subject: Re: propositions

[Martin Taylor 921113 16:00] (Bruce Nevin Fri 921113 12:34:38)

>>Pask argued that third-level feedback gave an iron-clad guarantee that
>>R had received what O intended to send. Such an agreement seems to me
>>to be more than a simple agreement about a simple perception, and far
>>more than 'an agreement like "yes, we employed that string of words."
>
>But (3) is not an agreement about a proposition, unless you mean
>"proposition" in some prelinguistic sense, e.g. a supposed mentalese.

Not "prelinguistic" but extralinguistic. I assume that the reference is to an imagined perception related to the feedback that agrees with the perception that gave rise to the original communication. If I can leave Pask and use LP theory, the Primal Message is a the reference state for the Originator's (O's) belief about the Recipient's (R's) belief. It has nothing to do with language.

The difference between the Primal Message and the current state of O's belief about R's belief engenders a virtual message (the error signal transformed by the output function). This virtual message is transformed through several such stages (Primal message - current belief -> virtual message) into a stream of words or other actions. The three Paskian levels of feedback correspond to an affirmation to O that R received the virtual message at that level, and that the proposition represented by the top-level Primal message is now understood by R. The top level of feedback (which could be well above the Paskian three) provides the input to the perceptual function of the ECS (protocol node) that output the top level virtual message. If O now believes that R's belief matches the reference, the message has been understood. The proposition was received. Nothing necessarily linguistic about it, unless O used language at lower levels to transmit the message. It's neither prelinguistic nor nonlinguistic.

>If we are interposing some third sort of entity between language
>perceptions and nonverbal perceptions, please elaborate.

Do you think we are? I don't, and I am not really clear about any distinction between "language perceptions" and "non-verbal perceptions," other than that a certain degree of consensual encodingism is permitted in language--we call it "syntax," and possibly "semantics." So I might be thinking of a second entity, though I doubt it. "Mentalese" is the language of control systems. It refers to the CEVs that correspond to perceptual signals. Words are part of mentalese exactly as are windows or social relationships. Signals flow.

Martin

Date: Fri Nov 13, 1992 4:50 pm PST
Subject: Lousy argument

[From Dag Forssell (921113)] Unsigned, but I presume Avery:

>Surprised not to have heard much for forsell & folson -- maybe my little
>blasts are having some effect. Almost reach the flame level today!

Please clarify. I put out a major post on Promoting PCT (921108) and later some related comment: Study PCT first. Then make revolution. (921109). You also allude to Hank Folson on language (921109).

Bill supposedly hit us over the head with a two by four, but did not cite a single specific of anything we had misunderstood or misstated.

You said: >Avery Andrews 191112.1013

>Recently on csgnet I've seen instances of a really lousy argument, to
>whit, X doesn't understand PCT because they're controlling for believing
>something else....

I have not said this quite this way. If you think so, please quote me in context. If I said something close, tell me whether it is true or not. From my perspective, my posts have been met with almost total silence (indifference?) and that is fine. If you think I have said something which is not true, please point it out. I will be quick to change when I see a better way to understand and express that understanding.

>Martin Taylor 921110 10:00

>>PCT shows clearly that people attend to that which they are concerned
>>about.

>

>That's one claim I've never heard on CSG-L before, and I don't believe it
>to be true. Most of the perceptions that are under good control are not in
>attention, and I'm not at all sure that "concern" is an attribute of an
>Elementary Control System. Where emotions and such feelings come into HPCT
>is a very interesting and hardly touched area.

I am not addressing Elementary Control Systems. I used the word concern to indicate a family of error signals in the upper reaches of the hierarchy. If you study biology, you will have a lot of error signals as to what you can't yet explain, and you will work on that.

I have tried to point out the difference in emphasis that results when A) Your primary interest and concern is in studying a given field in the life sciences and you try to use PCT as a tool to understand, or (What Avery alludes to) verify what you already know. And B) Your primary interest is in understanding PCT and you have no preconceived notions whatsoever.

I believe I have accurately portrayed some real phenomena. Please show me the error of my ways.

Best, Dag

Date: Fri Nov 13, 1992 5:17 pm PST
Subject: Re: PCT slogans

[Martin Taylor 921113 17:10]

Rick Marken wants a recap of the Vallacher and Wegner paper on "What do people think they are doing." Here are the relevant postings from Mark Olsen, Rick, and me.

From m-olson%UIUC.EDU@VM1.MCGILL.CA,,Feb,28,,1992
From: m-olson%UIUC.EDU@VM1.MCGILL.CA
Date: 28 Feb 1992
Subject: hierarchy attention and closed loop

To all, I just finished reading an article by Vallacher and Wegner, "What Do People Think They are Doing? Action Identification and Human Behavior," in Psych Rev 1987, 3-15. Its obviously been out for a while so maybe many of you already know about it, but nevertheless I thought I'd mention it. It concerns how individuals describe what they are doing, noting how they may focus on the goal state of any one of various levels of (the) a hierarchy. In one particular experiment, subjects were given coffee cups, one normal and one quite irregular. When asked later to describe what they were doing, descriptions

concerning the normal coffee cup focused on higher levels (getting my caffeine fix) while descriptions concerning the irregular cup focused on lower levels (drinking, lifting a cup to my lips). The relationship between attention and location of error in the hierarchy is evident here and they point this out. Their theory has something to say about causal attributions and variability of self-concept. As far as I can tell, it fits right into CT--Powers is even cited. They suggest that somehow their model is better, but I don't see how or why. It's a quick read so I'll let you all read it for yourselves rather than summarize the whole thing.

...

Carpe' Diem

Mark Olson

From: marken%AERO.ORG@VM1.MCGILL.CA, , Mar, 4, , 1992
From: marken%AERO.ORG@VM1.MCGILL.CA
Date: 4 Mar 1992
Subject: Psych Review Article

[From Rick Marken (920304b)]

I forgot to thank Mark Olson for calling attention to the Psych Review article by Vallacher and Wegner and for his comments on it (it's nice to have you back Mark). I haven't gotten a hold of it yet but I plan to get it this weekend and look forward to some good laughs (er -- interesting reading).

> They suggest that
>somehow their model is better[than PCT], but I don't see how or why.

Well, maybe they are right. Boy, would THAT be a surprise.

Hasta Luego Rick

From: mmt, , Mar, 4, , 1992
From: mmt
Date: 4 Mar 1992 17:20:36
Subject: Re: Psych Review Article

[Martin Taylor 920304 17:00]
(Rick Marken 920304b + Mark Olson)

>
>I forgot to thank Mark Olson for calling attention to the Psych Review
>article by Vallacher and Wegner and for his comments on it (it's nice to
>have you back Mark). I haven't gotten a hold of it yet but I plan to get it
>this weekend and look forward to some good laughs (er -- interesting
>reading).

>
>> They suggest that
>>somehow their model is better[than PCT], but I don't see how or why.

>
>Well, maybe they are right. Boy, would THAT be a surprise.

>
I see no such claim in their paper. The only mention of Powers is in a list of people who have proposed hierarchic theories, unless I missed something. What they do claim is

moderately interesting and in no way incompatible with PCT, though not supportive either. It may bear on the old discussion of the function and place of consciousness.

A gross simplification of their claim is that if the actions that allow high-level control are easy, then what subjects see themselves as doing is what we would call satisfying the high-level reference. But if the lower level control structure is disturbed or not well structured (the actions are more difficult), then people see themselves as "doing" the low-level things. It's the difference between "visiting Aunt Dorothy" and "trying to find the doorbell (for Aunt Dorothy's door)" (Not their example). Their claim also is that the behaviours that the subjects identify themselves as doing are less susceptible to disturbance than are the other levels. If someone is "keeping fit" then rain may cause them to switch from jogging to indoor exercise, whereas if they are "jogging" then rain is less likely to stop them. (Like, but not the same as, one of their examples).

In PCT terms, the idea seems to translate as:

(1) If an ECS is maintaining good control, then it is not high in conscious awareness, and (2) the conscious identification of a reference signal serves to stabilize that reference signal. The reference signal itself seems to become a percept for another "conscious" control system.

In a sense, the reference must be a percept, because the person holding it describes himself as perceiving it to be one. But it is also functioning as a reference signal at the same time, in the normal way. It seems to me that this idea alters the standard wiring diagram in an interesting way that might relate to the issue of conflict and reorganization.

Thanks for pointing out the reference, Mark.

Martin

Date: Fri Nov 13, 1992 5:38 pm PST
Subject: Re: associative memory

[Martin Taylor 921113 17:15] (Bruce Nevin Wed 921111 10:12:08)

> Whatever the case may actually turn out to be, I
> have assumed that associative memory introduces (and functions in)
> pathways over and above those used by the familiar p, r, and e signals
> in our diagrams of an ECS.

It may be so, but I hope it is not. I'd like to see what we can do with the simplest structures before postulating totally new ones that smack of the antiquity of psychology (do you like oxymorons? Are opposed syllables in Japanese Oxymora?).

The strength of association is often determined by asking a person to say the first word that comes into the head when a priming word is presented. The experimenter says "Black" and the subject says "White." If this happens often enough, White is thought to be associated with Black. An intriguing fact about this sort of experiment is that children tend to do it differently. A child is more likely to say something like "Sheep" in response to "Black." The child's association results are said to be "syntagmatic" rather than "paradigmatic." In other words, the child is more likely to say something that works with the primer word, linguistically or otherwise, whereas the adult is more likely to say something that has a function like that of the primer word.

In earlier (pre PCT) writing, I have speculated that what happens is a kind of holographic mirror process that is well developed in the adult, but not in the child. In the child, "Black" might be associated with, say, Night, a Hat, Sheep, and a few other things. In their turn, Night might be associated with a star, Hat with Daddy, Sheep with white and with grass, and so forth. I am, of course, talking about the child's experience with black things and things labelled "Black." In an adult, Black is associated with a vast multitude of perceptions, linguistic and non-linguistic, many of which are also associated with White, but few with Star, Daddy, or Grass. So for the adult, White is more strongly associated with Black than are any of the first-level "syntagmatic" associations.

From a PCT viewpoint, this transition seems interesting, if we notice that the kinds of sensory input to an ECS controlling for Black (word or visual percept) are very like those for an ECS controlling for White (indeed, they may be the same ECS, using different levels of a reference signal). But that cannot be the case for Black and Sheep. In the case of Black and Sheep, both together are inputs to the same perceptual function that can contribute to the "black sheep" percept (I don't want to address whether this is a scalar or a complex percept in this posting...it is irrelevant here).

One may assume that most consciously noted perceptions are disturbances. They represent change. The conjoining of Black with Sheep is a particular kind of disturbance and results frequently in a particular kind of error, and in a reasonably stable world in a particular kind of action, such as saying "black sheep." If there are not many other "black" experiences (and I am talking more about salience than about frequency here), the black->sheep association will tend to be observed in experiments because the "Black" part of the input drives the "black sheep" function more than it drives any other that results in a verbal action as a consequence of the disturbance.

With an adult, the same "black sheep" function may remain, but there are many other "black-x" functions. There are also many "white-x" functions with the same values of x. Now, "Black" induces a disturbance in most of the "black-x" functions. They correct in imagination by imagining "x" for all these x values. Many of them evoke a disturbance in "white-x" functions that is corrected by imagining "white", and the overall effect on "white" is sufficient that it is stronger than any of the intermediate "x", given that there is a reference somewhere that requires that a word be spoken. That word would be "white" rather than any of the "x".

As an adjunct to the above, it seems probable that percepts become fractionated with age, so that where once there might be a "black sheep" and a "black night" but no "black", later the abstract "black" is fractionated out of the various complexes. If that happens, both the input and output connections of black (and "black") will be very similar to those of white (and "white").

Pure speculation, but it provides a way of looking at association without introducing new pathways, using standard ECSs connected in a standard way. How to test it? I don't know.

Martin

Date: Fri Nov 13, 1992 5:53 pm PST
Subject: Re: information

(penni sibun 921113.1300)

[From: Bruce Nevin (Fri 921113 09:52:14)]

Thanks for posting the quote from Alice. It saved me typing it! (You must have an on-line copy. Is that available someplace?)

we have in in a misc archive next to the constitution. sorry--don't know where it came from!

i don't think we're going to agree on what is meant by meaning, but i think i understand you well enough on that point not to bother to argue any more.

>i think you are, roughly, equating words w/ verbal perceptions.

There is nothing but perceptions available to us. The distinction I make between verbal and nonverbal perceptions has to do with convention, arbitrariness, and participation in constituting linguistic information.

i used the word ``equate'' because you were in your previous message seeming to imply that all and only words are verbal perceptions. in this text here, you seem to be leaving open the possibility that things other than words c/b verbal perc's. i think that since verbal perc's are distinguished for you, you need to be very clear about what is and is not one.

In a code, there is an external table of equivalences between elements of the code and elements of language. Linguistic information is in a coded message, modulo substitutions through the table, precisely to the extent that it is in the decoded message. (Perhaps there is some legitimate use of "code" that does not rest upon language, but I have yet to encounter it. I believe that any other use of the concept of encoding is in fact only metaphor.)

i don't get the point of this para. you seem to be saying that a binary core dump and dna are not codes?

Saliency in the associated perceptions (meanings) depends upon control of those perceptions on the occasion, not upon something like a grammar in which things are looked up and correlated.

this is an example of where what you are saying is not specific to language. when i look around my office, some things are ``salient'' and some are not.

I should think this view would be very congenial with yours, in at least some respects.

if you didn't insist on distinguishing verbal and nonv perc's, or if i understood (and accepted) why you did, then i think we'd have a lot of common ground.

To be sure, there are constraints. Language is a matter of learned convention (language-particular) within the general characteristics of perceptual control (language universals). The discussion with Avery about agreements, membership, consensual reality, and so on bears directly on these constraints. That discussion suggests how there can be stabilities without tables for encoding.

these constraints are not particular to language.

In the two social situations that you described, you should be aware that your perception of what you are doing is quite likely different from the perceptions that other participants have.

of course i'm aware of that: i'm an adult of at least average competence. i also realize that the exact same is true when language is involved. (i mean, i used the precise term ``equate'' and you completely missed it, replying instead w/ a standard pct line. geez. ;-))

In the first situation in particular you are deliberately and probably effectively dissembling. In the second situation, you and your flirtation partner are very much involved with imagined perceptions. Is it necessary to say that there are many misunderstandings in this process?

perfectly reasonable observations about party patter and flirting. they also apply to friendly conversation and email msgs.. you've yet to say something convincing (to me, of course!) that distinguishes verbal from nonverbal.

That it is only a beginning, an assertion of availability and interest, and that it goes nowhere without on subsequent occasions spending some quality time (Ed's definition of the term) together? We call this evocation of perceptions and negotiation of agreements information, of course, but that is an informal, metaphoric use of the word, not the technical sense of linguistic information that I am using.

in my description of flirting, i gave specific instances of information being conveyed (eg, ``i like it when you do that''). why did you rewrite my description as ``an assertion of availability and interest''? i don't really know what that means, but that's not what i was talking about. i was referring to deliberately (most of the time?) and conventionally (most of the time?) collaboratively communicating. arbitrarily much ``info'' can be exchanged or collaboratively created, w/o words; the test of this i would guess is that the participants ``know'' more about each other afterwards than before.

cheers. --penni

Date: Fri Nov 13, 1992 6:22 pm PST
Subject: Re: Language, Interactionism

(penni sibun921113.1600)

i need to seriously not type for a while, so i'm going to opt out of all threads unless someone wants a very specific brief answer. nothing personal!

[From Hank Folson (921110)]

I am ignorant of interactionism and what it is all about. Okay, so I haven't been reading your posts in detail. That is because I am not controlling to

learn a lot about linguistics. PCT says that if you are not interested in something (controlling for it, in PCT terms) you will not observe it, or it will go in one ear and out the other (to use my mother's term, as she often does with me).

the interactionism thread was a couple months ago and didn't touch on ling at all, as i remember. i'm really not so interesting in ling per se, but in how language fits into everything else, which means spending a fair amount of time on ``everything else.''

1. Will you give me a couple of sentences long grossly oversimplified outline of interactionism?

int. is a blanket term that covers work done by a variety of people in a variety of disciplines. int. is a theory of how organisms do whatever it is that they do. int. doesn't believe in/emphasize an inside/outside split bet. an organism's head and ``everything else.''

it also doesn't believe there is a single locus of control the understanding of which will lead to the understanding of the organism. interactionism focuses, as its name implies, on the interplay of any and all parts of the situation to be understood/described/explained; also, as the name implies, it focuses on ongoing activity, not a static snapshot.

2. Who developed interactionism?

3. About when did they develop it?

major threads include: heidegger (philosophy, early 20thc); gibson (psychology, mid 20thc); garfinkel (created ethnomethodology/conversation analysis out of dissatisfaction w/ sociology/anthropology, late 60s); suchman (ethnomethodology and artificial intelligence, 1987) and agre&chapman (a.i., 1987)--this thread has been called situated activity.

4. What did (s)he believe in before (s)he came up with interactionism?

i'm not a mind reader! seriously, that's a very hard question. (don't get me wrong: i think history of ideas is *very* important. but i'd want to read a whole biography, not just a synopsis of someone's life.) i do know the last two, and what i know about them is that they were a couple of very bright boys in a rich intellectual environment w/ personalities and intellects that caused them to regularly believe six impossible things before breakfast.

5. What did his/her old belief lack that interactionism offered?

generalizing across researchers and fields, noninteractionism concentrates on individuals of various sorts, which are typically taken to be given (eg, indiv humans, particular social institutions). people drawn to int. want to look at the whole picture, and the continuously evolving picture; they feel that the typical static platonic way of looking at things probably completely misses what the hell is really going on.

6. How well has interactionism worked for those who use it?

i don't think anyone who has adopted it has given it up. i think people who use it believe that it has satisfactory explanatory power (esp. as compared to the competition!) and it generates interesting hypotheses. but that something ``works' could mean a lot of things.

7. Has any independent research backed up interactionism's claims?

i don't understand this question.

cheers. --penni

Date: Fri Nov 13, 1992 7:27 pm PST

[Avery Andrews 921114.1216] (Dag Forssell (921113))

The anonymous posting was a mailer mishap, but also a stupid and fatuous remark, for which I apologize. What I was mostly thinking of at the actual time was not the 'lousy argument' form, but the modelling, science, engineering, etc. issue. E.g. why people seem to think that we linguists are perverse in our approach to modelling. I do think there is a serious issue there.

(the part about hopping was about a frog population in a local nature reserve).

Avery.Andrews@anu.edu.au

Date: Fri Nov 13, 1992 7:32 pm PST

Subject: fundamentals

At the rate things go on the net, this might be an oldie but goodie; nevertheless:

Bill Powers (921110.1430) --

>Unchangeable belief is the enemy of control.

Yes. A nice clarification of PCT slogan # 2.

> I don't think anyone changes a system concept without turmoil,
>pain, and confusion, and the result is not likely to be predictable.

I agree. But system concept references do change -- probably as a result of reorganizations aimed at controlling intrinsic variables. My point WRT Ed's claim (about forsaking system concepts for social acceptance) was just that it implies that system concept references can be varied to achieve other goals (in Ed's example, social acceptance). Thus, Ed is placing control of system concepts lower than control of social acceptance in a control hierarchy. That's fine -- the point (which, of course, you agree with) is simply that references for perceptions that can be varied in the service of controlling other perceptions (and Ed implied that this could be done with what he called "system concepts") MUST be able to be so or there will be no control.

>This is why PCT is slow to be appreciated as a new system concept. Much
>easier to see it as a couple of new principles, or a set of clever
>calculations.

This is a VERY juicy point -- worth much more discussion, I think. PCT is not appreciated as a new system concept until one sees that it is really about themselves -- and their own experience; it is about life -- the only one that is really experienced from the point of view of the model itself -- one's own.

Have a nice weekend. Rick

Date: Fri Nov 13, 1992 7:37 pm PST
Subject: pasking levels

[Avery Andrews 921114.1250]

Surely people rarely in real life do anything half as hard as controlling perceptions at the Pasking levels. Mostly people control for the other person looking as if they were paying attention & understanding. When this does not seem to be happening, they do something to attract attention (grad students who I'm familiar with might say 'hey, did you get that', or something), but there are very strong social constraints on serious checking of comprehension.

For example, one of the things I discovered how to do in school was to access the contents of the short-term buffer that holds what someone is saying, even if I wasn't paying any attention to it. So the teacher could see that I wasn't paying attention to what he was saying, & ask me what he had just been saying, & I could repeat the last sentence more or less verbatim, thereby frustrating his attempt to conduct a public humiliation. But only certain sorts of people, such as teachers and wives, are allowed to even try this test.

My guess would be that teaching is typically so ineffective mostly because for ordinary communication (what will we have for dinner, how shall we organize the day, given 4 people, 6 commitments & one car), the 'second order' shortcuts to assessing what people know are entirely sufficient, & people have very little idea about what to do when they are not.

And of course a lot of talking is not an attempt to communicate any propositional info at all - Australians call a person who indulges in this a lot a 'motor mouth'. My guess is that the linguistic machinery serves so many masters that it is quite capable of rattling on all by itself. E.g. the reference signal is 'be saying something', & stuff comes out subject to some constraints, but not chosen for any particular reason. Under these circumstances I don't see how there can be anything like a 'Primal Message'.

Re Hamming distances, etc., I may have missed the relevant bits in the postings, but is there any place where the fundamental ideas are layed out in book or article form? Does it have anything to do with the theory of information & communication associated with people like Gardenfors (trying to connect it to a book I have a copy of). From my very likely faulty grasp of the postings, I see this as a regime where the errors really can go to zero, & stay there, because over shortish time-scales, people are expected to remember what you tell them.

Avery.Andrews@anu.edu.au

Date: Fri Nov 13, 1992 9:17 pm PST
Subject: Re: Hidden agendas

From Bill Cunningham (921113.2050) Re Bruce Nevin (921113.0953)

Sorry, I'm at home and can't download & edit.

Agree with you comments up to sentence beginning "The difficulty of making propositions binding..." Our purpose is to maximize the channel, as constituted by the (individually held) perceptions about the mutually held proposition set. Both parties may be considered as controlling for this perception. Yeh, there is error, but they are both working to reduce the error as seen by each. But that action is done in the present and does not require a performative. Our whole concern (at least for now) is for the maintenance of the channel at the time of information transfer. The notion of making the propositions binding isn't really something we require. (We can datestamp them if necessary, since perceived truth at a given time (past) is part of our process.

Child's play is the wrong term, in my view. I would agree heartily with "simple" or "fundamental". I would sign up for "trivial", except that it's a lesson that most of us relearn all too frequently. That's what triggered the remark by John Gabriel in the first place. I had made the (obvious) comment in the abstract and he suddenly applied it to a personal situation. But child's play it is not--because children simply don't play that way. Society (generally) views expansion of the mutual proposition set as a sign of maturity. We refer to people who refuse the attempt as "childish."

Bill C.

Date: Fri Nov 13, 1992 9:47 pm PST
Subject: Millis Outcomes List, Lousy Argument

[From Hank Folson (921113)] Bruce Nevin-

This is a very good list (Thu 921112 09:15:48), and I like its good intentions. This sort of list is usually used to direct very specifically focussed threats and force at parents, administrators, teachers and students. PCT says this will be very effective. We can say there is no need for further involvement, other than to offer congratulations. But, if you feel uncomfortable with this, apply The Test to see if the Millis parents are controlling not just for WHAT their children are taught, but also for HOW they are taught. Only then can the more positive options of PCT be offered with some hope of them being accepted by the Millis parents.

Avery Andrews (191112.1013)

>Recently on csgnet I've seen instances of a really lousy argument, to
>whit, X doesn't understand PCT because they're controlling for
>believing something else.

This is a very good point, and I agree with what you say. There are many things that X and I understand, but do not believe, because our beliefs lie elsewhere.

However, if Y is both controlling for some other theory and also controlling NOT to believe in PCT, Y can not increase his understanding of PCT.

In a post from Penni Sibun (921109.1630), I understood her to be closer to a Y, but having glanced at her post to me today, she may be closer to an X. We won't know until her injuries settle down which she thinks she is. Take it easy, Penni.

Best wishes, Hank

Date: Sat Nov 14, 1992 12:16 am PST
Subject: engineering, science

[From Dag Forssell (921114.0001)]

Avery Andrews 921114.1216

>....What I was mostly thinking of at the actual time was not the 'lousy
>argument' form, but the modelling, science, engineering, etc. issue.
>E.g. why people seem to think that we linguists are perverse in our
>approach to modelling. I do think there is a serious issue there.

Thank you, Avery. I too think there is a serious issue there somewhere.

The piece I posted Sunday night was a draft of an attempt to clarify to myself and potential customers what is so special and powerful about PCT. Christine and I had a good experience presenting our three days of training to a group of engineers. Afterwards, I was told by the HR person, that my program would not do for the manufacturing people. I was not dynamic enough. I did not give them nearly enough tools to take with them that they could use right away. My conclusion is that this person who apparently is very capable as a counselor and comfortable with conflict, did not listen to our program for the theory, but for the dynamics "(I know what people want!)" and for the values we expressed in the applications. This HR person majored in math once upon a time. The engineering people loved the program. They understood that it helped them understand themselves and their interactions.

I am an engineer by chance of my upbringing and all that. I am very well aware that my wife many times (and others with her) consider me head-in-the- clouds impossibly theoretical.

When Christine involved herself in sales in 1976, I went with her to some seminars and discovered motivational speakers. I wondered why some people bought and some did not. I started reading and followed a long trail which eventually brought me to PCT.

As you know, my interest in PCT now is to teach it and show people how to apply it. I am convinced it gives us a valid explanation of how we function as human beings. With open questions it is better than anything else by far.

I do not in any way mean to toot my own horn or put anyone else down. We are not talking human value here. My oldest daughter majored in linguistics at Stanford and teaches German now in Palo Alto.

But I do think that it makes a difference how you have been trained to think when it comes to understanding PCT. I think it is easier for engineers. If that is so, and I am one, common sense suggests that I concentrate on engineers as I promote myself and PCT, then let them promote PCT to others. My piece reflects such a marketing bias. I shall present our program as especially suited for engineers. I should have said this as I posted.

Why do I think it is easier for engineers? (Please note: It is only a matter of degree). I have learned to visualize stresses in materials according to specific explanations, compute and predict. I have also built and tested things which bent and broke just exactly the way I thought they would. I have learned to visualize and predict flows of fluids and electrons. It seems to me that somehow, the engineering training gives you a CONCEPTUAL LANGUAGE, a way of thinking and visualizing (non-verbal!) which stays with you all the time. I catch myself often visualizing happenings in terms of forces between objects, dynamics of motion, pressures etc. I almost never see evidence of this in my non-engineering friends in the same way.

Bill Powers just wrote a beautiful piece on: Words, meanings, computations (921102.1230) where he explains about voltages varying in an analog computer. If a person is contemptuous of theory (I perceive a majority to be) or thinks of theory as description or identifies theory with math logic, how can that person visualize the voltages or rather neuron currents Bill writes about. How can that person develop an appreciation for PCT in any depth, which is (I believe) helpful for the rapid development of a deep conviction.

Please comment on my discussion (921108) of the four kinds of theory. Does it make sense? I really want to know, so I'll quote myself here:

There are many levels of scientific rigor and interpretations of the word theory.

To some, "theory" is a "four letter word" signifying head-in-the-clouds irrelevant hokus-pokus. All theory is incomprehensible and irrelevant. From this perspective: You know what they say about engineers.

To some, "theory" means a description of a phenomenon. It is hard to test descriptions of phenomena in a rigorous way. Words are defined by other words.

To some, "theory" means a logical structure built on well defined hypotheses, which stand apart from the real world. This kind of theory is subject to logical proof.

To some, "theory" means precise description, explanation, prediction and test. Testing is done in the "real world" of hardware and is subject to small errors due to things like manufacture and measurements.

A person can certainly be aware of different meanings of the word theory and use it differently in different contexts.

To me, theory is the information / understanding / expectation a person uses to guide activity, no matter what the person calls it or how clear / valid it is as a map of that REALITY. In this view, everyone is full of theories and nothing but. But the quality varies.

The fourth kind is the engineering/physics approach. I am comfortable calling PCT an engineering science. I believe it is. I think that makes a sensible distinction. I think that nothing short of an engineering approach to theory is adequate to understand it. Mind you, a kid working on a car applies an engineering theory approach. I am not talking about degree snobbery here, but a way of thinking.

I believe some life scientists in some areas (medicine, biology, zoology, botany???) are very much into engineering thinking, others may not be. Linguists happen to be on this net, so I ask you: What is your definition of theory, what are your requirements. Do different linguists use different criteria.

My linguist daughter has not studied the engineering sciences much. She also does not worry about neuron currents, but is comfortable with my interest in PCT. She has understood the basics to her satisfaction and she tolerates dad. She is very capable and I love her dearly.

>E.g. why people seem to think that we linguists are perverse in our

>approach to modelling.

Having written all this, it occurs to me to ask you: Is your approach to modeling along the lines of math logic or is it along Bill's lines, a description of a physical structure and physical happenings that will actually do something. A different definition of theory and different criteria for testing and validity go with each. This may be an illuminating question.

Best, Dag

Date: Sat Nov 14, 1992 5:13 am PST
Subject: Text on-line

From Bill Silvert [921114] >(penni sibun 921113.1300)

>we have in in a misc archive next to the constitution. sorry--don't
>know where it came from!

Since quotes seem relevant to this group, it should be mentioned that vast quantities of literature are available in electronic form and can easily be obtained. Although I don't have the Internet information, there is a Project Gutenberg to enter virtually all literature into databases that can be freely accessed.

I recently purchased a CD-ROM that has about 600 Mbytes of literature for a small amount of money (\$25 or \$40 US). It contains the full text of Lewis Carroll's writings, several versions of the bible, all of Shakespeare's works, the complete stories of Arthur Conan Doyle, the Koran, etc. etc. etc. One page of text is about 2K so 600M is about 300,000 pages, a substantial library.

Bill Silvert

=====

Date: Sat Nov 14, 1992 8:03 am PST
Subject: language agreements

[From Bill Powers (921114.0710)]

Bruce Nevin (921112.1039) --

>In this we (you and I right now) are of course negotiating
>agreements and consensus. Has such an agreement any status apart
>from an individual reader's private perceptions? I suggest that
>there are prior agreements that it does, such as conversational
>implicatures and the conventions by which an utterance may
>constitute an instance of a speech act, and I believe that you and
>I both participate authentically in such agreements.

There are agreements that it does, and they are quite authentic -- in my perceptions. I don't know about yours.

A lot depends on what we are trying to talk about. If we simply want to USE our brains to talk about something else, like the surface manifestations of language, then we can assume that we all agree to the same conventions which we all perceive in the same way, and that my language structure is identical to yours. Such errors in this assumption as

show up in the course of interactions can be ignored or further negotiated until we each perceive no error.

If we want to build a brain model, however, we have to switch to a different view of language in which language is something that goes on exclusively in individual brains, not in the space between them. To describe what goes on in the space between them we use the physics model, which contains no agreements, but only properties and laws. The physics-model can convey sound waves and visual marks from one person's outputs to another's inputs, but that is all. Agreements occur in brains. It no longer matters much what language is being used to talk about, or which meanings are attached to which words. What matters then is to model the way in which language can be about things, what the things are that it is about, and how language and meaning are connected. If we discern rules in the way language is used, the problem is not so much to explicate the rules as to understand what a rule is and how it can influence anything.

If, as in Harris' operator grammar, we see evidence that some words seem to require others to follow them singly or in pairs, the problem is to remove this requirement from the words where it appears to be and put it into a brain function where it must be modeled. What sort of brain function would be needed in order to treat some words as operators and others as arguments, and to create an apparent dependency of an argument on an operator? Once we have seen a few examples of this phenomenon, as brain modelers there is little point in our looking for more examples until we have figured out what sorts of processes the brain must carry out in order to give rise to the surface appearance that "operators" "take" "arguments." If we look at the words to find an explanation we will be looking at the wrong things. Words are just arbitrary and interchangeable markings. In themselves they are just configurations of no particular importance. Their significance lies in the systems are are using them.

Here we can only fumble our way forward, fitting the phenomena into the model and looking for ways to reshape the model to make the fit better. The operator-argument phenomenon would seem to belong at the sequence level in the model as it stands. Computational phenomena would seem to belong at the program level. Simple naming of objects would seem to belong at the relationship level, or the category level. There are probably aspects of the Chomskyian approach that belong at the principle level. What we call "agreements" may also belong at the principle level, to the extent that the agreements amount to generalities. Much of the machinery of language production belongs at the lower levels.

I think that you and I have reached a lot of agreements on these matters.

And of course all these definitions of levels may well have to be changed in order to accomodate the kinds of perceptions we find being controlled, particularly the orphans.

There is probably more than one way to create the observed phenomenon of language. It's possible, for example, that Avery's programs and Penni's programs are both possibilities that would work. It's possible that they would both work better if they incorporated Harris' operator grammar. It's possible that different people and different cultures use different principles, all producing language that some others can recognize and that can be made to make sense to those others. After all, the purpose of language is to evoke meanings in others in some systematic way, a phenomenon that we call communication. It's not to do it in some particular way. If meanings do appear to be communicated to everyone's satisfaction, what does it matter how this is done? If I use one convention to indicate subject actions on an object, and you use a different one, we can still learn each other to understand. We all understood Yoda. The rules are not so rigid that we can't vary them a little while still having communication go on.

Maybe I'm proposing that in the end, Chomsky will turn out to have been right: language is innate, and we are born with the ability to learn and use it. But with the PCT slant on it, the significance of Chomsky's proposal may be quite different. It isn't that there is some specialized inherited ability to produce language, but that language is simply one way to use our control systems for a particular purpose. Anyone who can operate properly at all the necessary levels of the hierarchy can make and use language. If you're born deaf and dumb, you can still learn to get your meanings across -- to evoke in others experiences sufficiently like your own to create a sense of a shared reality. We can make up rules of language just as we can make up rules for a game; we can learn someone's rules of language and eventually apply them easily and automatically just as we can do for tennis or poker or computer programming.

Well, excuse me for meandering, but your post started a train of thought and I had to see how it would come out. I think that the final idea is kind of interesting, at least at the moment. And now that I think of it, I think that Hank Folson said the same thing a couple of days ago.

Martin Taylor (921113.1115) --

I see that we have the same opinion of Gordon Pask. I suspect that he is probably the brightest person I have ever met. Unfortunately he has elected to avoid communication by cloaking his ideas in opaque language (you and I independently chose the same word) and almost deliberately avoiding being understood. I don't know how he is with people he trusts. What I learned from Pask is that being bright isn't enough. I also got the impression, the last time I saw him talking at a cybernetics meeting, that he had enormous contempt for the people he was talking to. He seemed to be saying "I know you didn't understand that, so here's something even harder for you to understand."

>Pask argued that third-level feedback gave an iron-clad guarantee
>that R had received what O intended to send. Such an agreement
>seems to me to be more than a simple agreement about a simple
>perception, and far more than 'an agreement like "yes, we employed
>that string of words."'

As Bruce has been pointing out, that "iron-clad" agreement rests on prior agreements: locus, distant and equidistant, points, plane. Somewhere it all comes back to demonstration. And demonstration rests ultimately on private perception.

Isaac Kurtzer (921113.1200) --
RE: cats raised under 8/sec strobe light.

What a beautiful example! I think this is very strong evidence that the transition level is ABOVE the configuration level. There seems to be some sort of transition level below as well, but it may be simply part of linking the control systems to a physical world in which acceleration leads to velocity which leads to position.

Latin quotes: not bored academicians: academicians trying to be sure that nobody mistakes them for normal people.

Best to all Bill P.

Date: Sat Nov 14, 1992 10:05 am PST
Subject: CATS: from Mary Powers

[from Mary Powers 9211.14]

David Goldstein:

This took a while, because someone else had the library book. From Catwatching, by Desmond Morris: "why does a cat lick its face when it is not dirty?"

...a telltale sign that a cat is becoming agitated, while at the same time being fascinated or puzzled by something ... "displacement activities" ... some animals nibble their paws ...birds wipe their beaks on a branch ... chimpanzees scratch their arms or their chins ... for felines the tongue swipe is the favorite action.

...cats do not like vibrating noises with a high pitch, but they are intrigued ... a coin rubbed back and forth along the teeth of a comb produces such a noise. Almost every cat ... stares at the comb ... and after a few seconds starts licking its lips ... this works for full grown lions as well as for small tabbies. [Sometimes followed by a sneeze or a yawn]

Why ... is something of a mystery [blah blah about evolution and rattlesnakes]

Mary P.

Date: Sat Nov 14, 1992 10:26 am PST
Subject: setting standards

from Ed Ford (921114:1115) (from Bill Powers 921113.1000)

>...we received your cards on Teaching Responsible Thinking. I think
>you need to take a closer look at what you're recommending, or at
>least the way you put it.

First, for those who are unfamiliar with my card, here it is...

TEACHING RESPONSIBLE THINKING

Responsibility - the willingness and ability of people to follow standards and rules and ultimately to set their own, without infringing on the rights of others.

A. ESTABLISH STRONG RELATIONSHIPS

The most important step when teaching responsible thinking involves spending daily quality time alone with each person. (explained on reverse side of card). People must believe two things: you care about them and you have confidence in their ability to resolve their problems. You will then be more effective when teaching responsible thinking.

B. SET STANDARDS; ASK FOR CHOICES & CONSEQUENCES

You have to set specific and reasonable rules and standards that must be consistently applied over time and enforced fairly with each person. When they are not willing to follow standards or obey rules, ask them to name the various choices they have and to explain the consequences

that result from making those choices. The consequences should include the loss of the privilege which is related to the responsible choice they refuse to make.

Loss of privileges or restrictions must be the result of their not being willing to work at resolving their problem. Lifting restrictions or the return of a privilege should be tied to their having made a commitment and a specific plan to resolve their problem. Remember, trying to control people by rewarding or punishing them doesn't teach responsible thinking.

C. TEACH RESPONSIBLE THINKING

Ask them what they think - keep questioning them; don't tell them what you think; ignore excuses & don't ask why; be non-judgemental; be specific & stay focused.

1. Exploration: Basic questions to ask are...

What is it that you want? What are the rules or standards?
What were you doing to get what you want?

2. Evaluation: Getting them to think responsibly...

Is what you're doing getting you what you want?
Is what you're doing against the rules?

3. Commitment: Getting them to choose responsibility...

Are you willing to work at resolving your problems?

4. Teach them how to create a plan

- a. establish specific area of needed improvement
(keep plan small to assure success)
- b. set a measurable goal for needed feedback
- c. have them think through then explain in detail how they're going to achieve their measurable goal
- d. set up a chart which shows progress in time increments (hourly, daily, or weekly, etc.)

picture of a chart

It seems to me that when two or more living control systems find themselves in the same environment, in order for them to live in harmony and cooperatively, they have to agree on a way things ought to be, a system of concepts, which are best expressed and set forth by agreeing to a set of standards upon which they base their choices as they attempt to find satisfaction while living together. (I see standards as synonymous with rules, criteria, principles, guidelines, etc.) Thus the needed harmony between levels of the hierarchy in social groups. As they live their lives, trying to satisfy their own individual goals, the choices they make, if based on agreed to standards, will more than likely make it easier for them to live in harmony with each other.

In the order of nature, we first learn to follow standards as children at home and then, ultimately, to set our own. For us to live in harmony, we must always set rules while respecting the rights of others. Whether at home, at school, at work, in an institutional setting, or just buying gasoline, we are constantly surrounded by standards and rules. Thus the need to learn to follow standards as well as to set our own standards while resolving our internal conflicts.

Any time I deal with anyone, whether in private practice or elsewhere, standards and rules are a part of life. Whether it is setting standards for the kind of spouse we want (thus to help us make a choice), or wanting to get along with a parent, or getting through school, or interacting properly on the CSGnet by how to sign on, etc., standards are a part of life. HOWEVER, IN ORDER TO HELP LIVING CONTROL SYSTEMS RESOLVE THEIR OWN INTERNAL CONFLICTS AND TO TEACH THEM HOW TO DEAL WITH THEIR LIVES, THERE HAS TO BE A BASIC UNDERSTANDING OF STANDARDS AND RULES AND OF HOW CONSEQUENCES AND CHOICES ARE INTEGRATED INTO THE STANDARDS CONCEPT ALREADY ESTABLISHED IN THE SETTING WHERE THEY ARE BEING TAUGHT. That is why I set forth a general understanding and the need for setting standards as well as establishing a model for where and how choices and consequences are integrated into the standard concept prior to explaining the style and techniques for teaching responsible thinking (in Section C) to anyone, any where, whether at home, at school, or where ever.

I think, Bill, you confused what I was saying in Section B as a way to work with others. I was merely setting forth the essential elements that go into helping people learn to discipline themselves and deal with the rules in environments where they find themselves or help those who have to establish rules and standards in a setting and how all the concepts interrelate. The approach I use when dealing with others, as set forth in the role plays and explanations in Chapters Nine and Ten in Freedom From Stress has not been abandoned. The attempt to teach responsible thinking through asking questions is written out in Section C. What I was explaining in Section B was the process, the standards to be considered and incorporated when teaching others how to obey rules or helping them learn how to function effectively in an environment where rules exist.

I use this card when teaching parents or when teach in residential treatment centers or schools or in businesses. The essential elements explained in Section B are the same. The style outlined in Section C, as you rightly recognized, is very different, depending on the person and the situation. I use the ideas for setting standards in Section B with my grandchildren as I do with juveniles in a detention center. But the style for working with people outlined in Section C can be very, very different.

>Maybe you should think about approaching people in institutions as you
>would people outside them. Teach them how to use PCT in the real
>world they'll be going back to, not the artificial and cramped world
>they're in at the moment.

As you can see from the above, I do, always have and always will.

Finally, my experience over the years has taught me that there are tremendous differences in the understanding of the role of standards, the meaning and place of consequences and choices. I was trying to establish as clear an explanation as I could of the whole concept. I trust my card makes a little more sense.

Ed Ford

Date: Sat Nov 14, 1992 2:13 pm PST
Subject: primer part II

Continuing with the draft version of the modeling primer:

Experimenting with the control paradigm
A primer for computer modelers

PART II: Finite gain and leaky integrators.

The model in Part I used a pure integrating output. In real physical systems, especially nervous systems, integrators are not perfect; they leak. The greater the output of the integrator, the faster it leaks. The result is to create the equivalent of an amplifier that multiplies the input by some finite number, but takes a while to follow changes in the input. For a step input, the output begins to rise as if a pure integrator were present, but instead of the output continuing to rise, it levels off at some multiple of the input magnitude.

We will look now at the control system of Part I with a leaky integrator for an output function. The program now is slightly more complex, in that we will plot the values of the variables against time, using a text-mode output on a graph that is 80 characters wide and 25 high.

This subject is discussed at some length because it has important implications for simulating continuous physical systems on a digital computer. After we have finished this section we will be ready to look at some new quantitative relationships found in control systems.

New program:

```
/* primer2.c */

#include "stdio.h"
#include "conio.h"

void main()
{
float sp = 0.0, /* initialize signals and quantities */
    sr = 20.0,
    se = 0.0,
    qo = 0.0,
    qi = 0.0,
    qd = 0.0;
float kd = 1.0, /* set constants */
    ki = 1.0,
    ko = 100.0,
    kf = 1.0,
    dt = 0.1
    ks = 50.0; /* slowing factor */
int i;
clrscr(); /* alternative: for(i=0;i<25;++i) putchar(0x0d); */
for(i=1;i<=80;++i)
{
    if(i > 40) qd = 9.0; else qd = 0.0;
    qi = kf*qo + kd*qd;
    sp = ki*qi;
    se = sr - sp;    qo = qo + (ko*se*dt - qo)/ks;
    gotoxy(i,24 - sr/2.0 - 0.5); putchar('r');
    gotoxy(i,24 - qd/2.0 - 0.5); putchar('d');
    gotoxy(i,24 - sp/2.0 - 0.5); putchar('p');
}
(void) getch(); /* wait for keystroke */
```

}

Notice the changes. We now iterate 80 times, with the index i starting at 1. The disturbance q_d is zero until the 40th iteration when it rises to 9.0. It remains at 9.0 for the rest of the run. With $dt = 0.1$ sec, the horizontal dimension corresponds to 8 sec of real time.

The slowing factor

A leaky integrator could be made from the integrator of part I:

$$q_o = q_o + k_o * s_e * dt$$

Simply by subtracting an amount proportional to q_o on every integration:

$$q_o = q_o + k_o * s_e * dt - q_o * dt * (\text{LeakFactor}).$$

There is another form equivalent to this that makes it possible to adjust the steady-state amplification factor and time constant independently. We will use this alternate form for the output function, as follows:

$$q_o = q_o + (k_o * s_e - q_o) * dt / k_s; \quad /* \text{ dt}/k_s = \text{slowing factor} */$$

To understand this function, first consider just the part in the parentheses. $k_o * s_e$ is the value that the output q_o would have if the error signal s_e were simply multiplied by an amplification factor k_o . Subtracting q_o from this value calculates how far from the current value of q_o the final value of the output is: this is the change necessary to reach a final state of q_o defined as $k_o * s_e$. If we just added $(k_o * s_e - q_o)$ to q_o , the result would be $k_o * s_e$ -- the value of q_o on the left would simply be equal to $k_o * s_e$.

This amount of change, however, is multiplied by dt/k_s . If k_s , a "slowing factor," is set equal to dt , then the entire difference between the current value of q_o and the final value is added to q_o , so that q_o (on the left of the = signal) becomes equal to the final value on the first iteration. If $k_o * s_e$ does not change, on the next iteration q_o will be equal to $k_o * s_e$, so the difference $k_o * s_e - q_o$ will be zero and nothing more will be added to q_o during later iterations.

On the other hand, if k_s is made larger than dt , so dt/k_s is less than 1, only part of the difference will be added on each iteration and q_o will approach the value $k_o * s_e$ in a series of diminishing steps. If, for example, k_s were made equal to $2 * dt$, then only half the difference would be added on each iteration. The output q_o would rise toward the limiting value $k_o * s_e$ by going half the distance, then half the remaining distance, and so on, on each successive iteration. With a one-unit final value, the steps would be $1/2$, $3/4$, $7/8$, $15/16$... and so on. To a first approximation, then, the slowing factor k_s sets the time constant of the output function. If k_s is set to 1.0, the time constant will be about 10 iterations, because dt is 0.1. In 10 iterations, the output will go about $2/3$ of the way to the final value after a step- change in s_e .

Loop gain and the optimum slowing factor

In this closed-loop system, the error signal does not remain constant on successive iterations. Thus we do not see the actual time constant of the output function in the overall behavior of the control system. The control system as a whole will show a much shorter time constant than k_s would imply. In fact, the overall time constant can be

reduced to a single iteration of the program no matter how long is the time constant in the output function.

Let us first define the loop gain of this system. The loop gain is the product $-k_o*k_f*k_i$: the product of all multiplication constants encountered in one trip around the closed loop (starting anywhere). The negative sign is introduced by the comparator, where an increase of 1 unit in the perceptual signal produces a change in the error signal of - 1 unit (reference signal constant). For calculating loop gain we are concerned only with the effects of small changes, not the absolute magnitudes of signals and quantities.

It can be shown that the optimum value of dt/k_s is simply $1/(1 - \text{loop gain})$, where the loop gain itself is always a negative number. With this value of dt/k_s , the control system will, after a step-disturbance, reach equilibrium on the first iteration of the program. Conversely, for any value of k_s there is a loop gain that will give the same result.

The above program is set up with a loop gain of -100 ($k_o = 100.0$, $k_f = 1.0$, $k_i = 1.0$). The optimum value of dt/k_s would be $1/101$. With dt set to 0.1, the optimum k_s is 10.1. The program is initialized with $k_s = 50$, or about 5 times the optimum value. This should produce a time constant of the whole control system of about 5 times the optimum value for reaching equilibrium in one iteration; it should therefore take the control system about 5 iterations to reach 2/3 of the final value. You can compile the program and run it now to see that this is the case. Notice that with an inherent time constant of 50 iterations in the output function, the overall system has a time constant of only 5 iterations.

On the plot, r indicates the reference value, p the perceptual signal, and d the disturbance.

Now reduce k_s to the optimum value of 10.1, recompile and run, and verify that the final value is reached after the first iteration.

Note what happens when the disturbance turns on (where the row of d 's suddenly jumps upward). The perceptual signal is disturbed upward, above the reference signal, for one iteration. Then it returns to the reference value and stays there even though the disturbance is still present. If you increase the value of k_s , you will see that the error correction becomes slower.

A digital artifact

If you make k_s smaller than dt , you will begin to see a computer artifact. For k_s in the range between $0.5*dt$ and dt , the approach to a final state will be oscillatory. If k_s is even smaller, the oscillations will begin to increase in amplitude exponentially; the system will run away. Try a value of $k_s = 5.5$ to see the oscillatory approach to a steady state. The nearer you get to 5.0, the longer the oscillations will persist. If you reduce k_s to or below 5.0, all you will see will be a few scattered p 's -- the growing oscillations quickly carry the trace off the screen.

This is an important phenomenon that shows the limitations of modeling a continuous system on a digital computer. The clue that tells you this is a computer artifact is simple: the oscillations (when they occur) always have a reversal period of one iteration of the program, regardless of the setting of k_s or dt . The constant dt expresses the meaning of one iteration in terms of real time. As long as k_s is greater than dt , the time constants of changes have physical meaning. But as soon as the oscillations begin,

they occur at a frequency of exactly one reversal per iteration, however you change ks. That frequency has no physical significance.

If you leave the value of ks at 5.0 and change dt to 0.01, you will now see a normal exponential rise of sp just as before. It won't be evident from the plot, but now the width of the plot represents only 0.8 seconds instead of 8.0 as before because one iteration now corresponds to 0.01 sec instead of 0.1 sec. The rise time of the p's is still about 5 iterations (to the 2/3 value), but now that corresponds to 0.05 sec or half an iteration with the old value of dt. This proves that the model system still behaves normally with small values of the time constant (or slowing factor), even though when the same model with the same physical parameters is run with a coarser time scale, it seems to become unstable. Control systems CAN become unstable. What we see here, however, is not instability in the physical system, but only in the digital representation of it.

The moral of this story is that when you employ slowing factors you must always make sure that ks is greater than or equal to dt. If you need to represent a system with a faster rise to asymptote, you must reduce dt. Then you can reduce the time constant further by reducing ks as required to match the behavior of the real system. We will see later how this kind of analysis works when the real system has some actual time delays in it.

Things to experiment with

Loop gain is affected by not only by ko, but by kf and ki. You can try varying these other constants to see the effect. You will always find that the fastest physically meaningful speed of error correction is a single iteration, and that to get this speed you must set kd/ks to $1/(1 + ko*kf*ki)$. If you want the system to have a time constant of x iterations, you simply multiply the demoninator by x. When you have created a certain time constant in the correction of errors due to disturbances, you can then verify that you get the same time constant in response to step-changes in the reference signal sr. To do this you can crib some lines from primer1.c in Part I. The reference signal value is already displayed. If you want to display the error signal, just add the line

```
gotoxy(i,24 - se/2.0 - 0.5); putch('e');
```

... in the appropriate place, and so on for other variables you would like to plot. Notice that when two or more variables plot to the same position, only the last one to be plotted can be seen.

Best to all, Bill P.

Date: Sat Nov 14, 1992 3:12 pm PST
From: Gary A. Cziko
Subject: Intro Doc References

Dag:

Thanks for reposting the Intro document. I see now what you mean about the format for the commands getting messed up.

For some strange reason I felt compelled to put the references in a format closer to that of the American Psychological Association. I have appended these to this message.

Thanks again for all your work on this. We are getting a steady trickle of new subscribers and I am very pleased to now have a good intro document to send them.--Gary

P.S. I just realized that it would probably be a good idea to send you back the entire document so that you can see how it will appear when I send it to the net. I will send this as a separate post to you. Let me know if it shows up all right.

+++++

LITERATURE REFERENCES

For a complete list of CSG-related publications, get the file biblio.pct from the fileserver as described above. Here are some selected books and papers on Perceptual Control Theory:

Powers, William T. (1973). *Behavior: The control of perception*. Hawthorne, NY: Aldine DeGruyter.

The basic text.

*Robertson, Richard J. and Powers, William T. (Eds.). (1990).

Introduction to modern psychology: The control theory view.

Gravel Switch, KY: CSG Book. (\$25 postpaid)

College-level text.

*Powers, William T. (1989). *Living control systems: Selected papers*.

Gravel Switch, KY: CSG Books. (\$16.50 postpaid)

Previously published papers, 1960-1988.

*Powers, William T. (1992). *Living control systems II : Selected papers*.

Gravel Switch, KY: CSG Books. (\$22 postpaid)

Previously unpublished papers, 1959-1990.

*Marken, Richard S. (1992). *Mind readings : Experimental studies of purpose*.

Gravel Switch, KY: CSG Books. (\$18 postpaid)

Research papers exploring control.

Marken, Richard S. (Ed.). (1990). *Purposeful Behavior : The control theory approach*.

American Behavioral Scientist, 34(1). (Thousand Oaks, CA:

Sage Publications)

11 articles on control theory.

Runkel, Philip J. (1990). *Casting nets and testing specimens*. New York:

Praeger.

When statistics are appropriate; when models are required.

Hershberger, Wayne. (Ed.). (1989). *Volitional action: Conation and control*

(Advances in Psychology No. 62). New York: North-Holland.

25 articles (not all PCT)

Ford, Edward E. (1989). *Freedom from stress*. Scottsdale AZ: Brandt

Publishing.

A self-help book. PCT in a counseling framework.

Gibbons, Hugh. (1990). *The death of Jeffrey Stapleton: Exploring the way*

lawyers think. Concord, NH: Franklin Pierce Law Center.

A text for law students using control theory.

McClelland, Kent. (1992). Perceptual control and sociological theory.
Not yet published. Available from the author,
Grinnell University, Grinnell, Iowa.

McPhail, Clark. (1990). The myth of the madding crowd. New York: Aldine
de Gruyter.

Introduces control theory as a way to explain group behavior.

McPhail, Clark., Powers, William T., & Tucker, Charles W. (1992).
Simulating individual and collective action in temporary gatherings.
Social Science Computer Review, 10(1), 1-28.

Computer simulation of control systems in groups.

Petrie, Hugh G. (1981). The dilemma of inquiry and learning. Chicago:
University of Chicago Press.

Introduces PCT to educational theory.

Richardson, George P. (1991). Feedback thought in social science and
systems theory. Philadelphia: University of Pennsylvania Press.

A review of systems thinking, including PCT.

*These items are available from CSG Books, 460 Black Lick Road,
Gravel Switch, KY, 40328.

Gary A. Cziko Telephone: (217) 333-8527
Educational Psychology FAX: (217) 244-7620
University of Illinois E-mail: g-cziko@uiuc.edu
1310 S. Sixth Street Radio: N9MJZ
210 Education Building
Champaign, Illinois 61820-6990
USA

Date: Sat Nov 14, 1992 3:16 pm PST

Subject: GA Evolution of Control Systems: Tracking the Bill Powers Trail

[from Gary Cziko 921114.2241 GMT]

I have been reading and enjoying Stephen Levy's Artificial Life: The Quest for New
Creation (New York: Pantheon, 1992), particularly his chapter on genetic algorithms
(GAs).

In this chapter he describes the work of David Jefferson and a group of UCLA researchers
who used GAs to "develop trail-following behavior in an insect made of information. The
task they selected, called Tracker, was following a specific trail of eighty-nine squares
on a thirty-two-by-thirty-two-square toroidal grid. . . . Inspired by actual pheromone
trails used by ants to aid each other in foraging, the trail twisted and turned, and
became increasingly difficult to follow as it progressed. It suffered gaps at several
points, and, by the last segment of the trail, there were more missing squares than
"scented" ones."

"It would be a mild challenge for any computer hacker to write a computer program for an
artificial creature to follow this trail, which the UCLA team nicknamed "The John Muir
Trail." But Jefferson and his colleagues hoped to get something from nothing--they were
counting on evolution to write the program." (p. 165).

Jefferson et al. were successful in evolving in their 70th generation a population of ants of which the majority could complete the trail perfectly. But, of course, these ants were living in a make-believe digital world where behavior was guaranteed to turn directly into consequences. No tilting landscape, no gusting wind, no obstacles to circumvent, no fatigue, no cascading of error in turning muscle accelerations into limb velocities and limb velocities into ant positions.

So why doesn't somebody out there do this honestly, that is, using GAS to deal with a computer environment in which disturbances have been programmed (or using a real physical device as the ant)? The trail would of course be nicknamed the Bill Powers Trail. We should then see a control system evolve. Although I am no programming whiz, this doesn't seem all that hard to do. Who knows, maybe one could even set up a challenging enough environment which would result in the evolution of hierarchical control systems.

Perhaps this would be a good dissertation project for some students out there.--Gary

Date: Sat Nov 14, 1992 4:11 pm PST
Subject: Interactionism & PCT

[From Hank Folson (921114)] (penni sibun921113.1600)

I can see why you are checking out CSGnet & PCT. Interactionism and PCT are looking at similar worlds.

>>1. Will you give me a couple of sentences long grossly oversimplified
>>outline of interactionism?
>interactionism is a theory of how organisms do whatever it is that they do.
>int. doesn't believe in/emphasize an inside/outside split between an
>organism's head and ``everything else.''

While PCTers often refer to this as a split or a dividing wall, etc., it can also be just as accurately described as an interface. Interactionism should have no problems with interfaces, as all organisms start and stop someplace, so they must have interfaces between themselves and what they interact with. Even if interactionists don't place much emphasis on it, there must be internal interfaces, too. For example, the body's sensors must interface with whatever the body's mental system is. This doesn't make PCT right, but I am saying that there must be an interface between our sensors and our brains for interactionists, too.

The way PCTers talk about ">the split between the brain and everything else" is aimed at S-R thinkers, not interactionists. Without backing off on what PCTers are claiming, the way to describe it to an interactionist might be: An organism's ability to interact is affected by the quality and capabilities of its sensors. An organism with limited sensors will interact less, and less effectively, than an organism with powerful sensors. A corollary is that an organism that has always had limited sensing capabilities will have no way of knowing that it is missing any interaction possibilities. (This argument is easier to buy if you apply it to organisms that do not have a better equipped organism communicating to them that they are missing out on something.) Any organism's interactions with the real world are limited by the capabilities of its sensors. Thus the organism intentionally interacts only with the part of the world that it can sense. This is where the PCTers start talking about the perceived versus external realities. The idea comes not from some basic philosophy of PCT, but from the unavoidable practical organic limitations of organisms that prevent them from perfectly sensing their environment. It

is a necessary evil, not a goal of PCT, to talk about different realities. The awareness of limited realities allows PCTers to better understand the interactions of organisms.

>it also doesn't believe there is a single locus of control the
>understanding of which will lead to the understanding of the organism.

Neither does PCT. Engineers stack simple control systems in a hierarchical fashion to achieve sophisticated control with minimum expenditure of time, resources and effort, so this is a well established possibility. Organisms use stacked levels of control systems for the same reasons, although I understand that more research must be done to make this a generally accepted truth. But if you accept on faith that organisms can do what engineers do with control systems, PCT claims several things:

1. Control is distributed, in that control signals need to go up the hierarchy only until no further errors signals remain. Simple inputs are evaluated and acted upon if needed at lower levels. Complex inputs must go higher up the hierarchy towards the brain.

2. If by "locus of control" you mean that control involves not just one organism, but also the organism(s) it is interacting with, then interactionism & PCT differ. PCT holds that all organisms are independent. The interactions between organisms are only related by the ongoing interaction of each controlling for what it wants, with the controlling actions of other organisms being disturbances to what each is controlling to maintain. So there is plenty of interaction, and it is purposeful, but it is coordinated in purpose only for those organisms that are controlling to coordinate their actions.

3. By understanding the number and nature of the levels, you can better understand how an organism works, and why it does what it does.

>interactionism focuses, as its name implies, on the interplay of any
>and all parts of the situation to be understood/described/explained;

I read your sentence as meaning that interactionists look at what is going on from the point of view of an external observer. If you want to pull Bill Power's chain, just say that that is what PCT theorists do. The PCT theorist will observe from the imagined position of being inside the organism using the organism's sensors only. The external observer will observe an organism bumping into things and have to choose from the possibilities that the organism is: aggressive, stupid, poorly coordinated, or has limited vision. The PCT observer will evaluate or model the organism's senses and know why it is controlling the way it is (although it is easier said than done).

>(interactionism) focuses on ongoing activity, not a static snapshot.

The definition of a control system is that it is continuously comparing inputs to reference signals. If control systems took snapshots, they would have big problems: One, they would need a good clock to tell it when to take the next snapshot. Two, the controlling would be very inefficient because there would be no reduction of error signals between snapshots, so the controlling action and effort required at each snapshot would be much greater, wasting valuable energy and yielding cruder control.

Thanks for answering my other questions. I'll just quickly take a stab in the dark and say that I guess that interactionists do a good job of describing WHAT organisms' interactions are, and a fairly good job of describing HOW they interact, but they come up short on explaining WHY they interact. PCT is not only all about how organisms interact, it is also about why they interact.

Interactively, Hank Folson

9211C

Subject: Judging Theories By Using PCT

[From Hank Folson (921115)]

Being interested in PCT applications, I came up with a list of 7 questions which are intended to be an application of PCT principles. The goal is to efficiently evaluate other theories, in case there is a better one than PCT. After all, there has to be something better than a theory that says: 1. You want to control everybody because you are an independent control system, and 2. You can't control anybody because they are all independent control systems. The world would be in a terrible mess if it ran on that basis.

Penni answered them about Interactionism, and here is my answering of them about PCT, with the reasons for the questions in []. Bill Powers can correct my sloppy scholarship.

>1. Will you give me a couple of sentences long grossly oversimplified
>outline of the theory?

See Mary Powers(921111) or Dag Forssell(921112)

[This just gives you a rough idea about what the theory covers. It may allow you to see that there is no relationship to your interests.]

>2. Who developed the theory? About when?

William T. Powers, Clarke and about 30-40 years ago.

[This gives you an idea about the background of the theorists and a time frame so you have some idea of how mature or complete the theory is.]

>3. What did (s)he believe in before (s)he came up with the theory?
Bill Powers was studying traditional psychology.

[This is straight PCT. The question is looking for the error signals that drove the new theory.]

>4. What did his/her old belief lack that the theory offered?

Traditional psychology is based on a stimulus-response model, which Bill felt did not fully explain behavior. PCT offered an explanation of how living organisms function.

[Again straight PCT. How did the new theory reduce the error signals in the creators?]

>5. Where did the basic idea of the theory come from?

Bill had been involved as a technician in WW II with control systems for naval guns. He made a connection between mechanical control systems and organically based ones.

[New ideas are the result of reorganization of our experiences, per PCT. Evaluating the meaning of this answer can be very difficult. What does aiming guns at airplanes have to do with how living organisms work?]

>6. How well has the theory worked for those who use it?

Well, it has given Bill Powers an interesting avocation... Ed Ford and others are successfully applying it to counseling and education. Franz Plooij has done impressive work in biology. And so on.

[This question gets to the applications of the theory. The proof is in the pudding.]

>7. Has any independent research backed up the theory's claims?

My two favorites are: Studies of E. coli bacteria by a non PCTer showed that it functions quite well using a very simple control system. This is an evolutionary stepping stone for the more complex control systems in humans. Plooij studied the nipple rooting instinct in monkeys from a PCT point of view. (In traditional psychology for "instinct" read "magic"). He found that the lactating monkey mommy has a temperature gradient across her breast increasing towards the nipple. The baby monkey has temperature sensors in its lips. A simple control system leads the baby monkey to move its mouth towards higher temperatures. When it bites down on the nipple, other control systems take over.

[This finds examples of real world proof or application of the theory for your evaluation.]

Note that if these questions were asked about PCT 30 years ago, the answers would not have been as complete. The questioner would have to wait for years for confirmation, or help to study/prove the theory, or decide for or against the theory with incomplete information.

Theoretically, Hank Folsom

Date: Sun Nov 15, 1992 9:53 am PST
Subject: Primer part III

Continuing with the draft version of the modeling primer:

Experimenting with the control paradigm
A primer for computer modelers

PART III: The effects of delays.

It is often said among the ill-informed that control systems will not work well in organisms because of reaction-time and other delays. In this section we will set up a control system with delays in it, first in the output function and then in the input function, to see what is required for such a system to maintain tight and stable control.

Output delay

The program is an extension of the one in Part II:

/* primer3.c */

```

#include "stdio.h"
#include "conio.h"

void main()
{
float sp = 0.0, /* initialize signals and quantities */
    sr = 20.0,
    se = 0.0,
    qo = 0.0,
    qi = 0.0,
    qd = 0.0;
float kd = 1.0, /* set constants */
    ki = 1.0,
    ko = 100.0,
    kf = 1.0,
    ks = 70.0,
    dt = 0.1;
float qo3 = 0.0, /* initialize variables for output delay */
    qo2 = 0.0,
    qo1 = 0.0;
int i;
clrscr(); /* alternative: for(i=0;i<25;++i) putchar(0x0d); */
for(i=1;i<=80;++i)
{
    if(i > 40) qd = 9.0; else qd = 0.0;
    qi = kf*qo + kd*qd;
    sp = ki*qi;
    se = sr - sp;
    qo = qo1;
    qo1 = qo2;
    qo2 = qo3;
    qo3 = qo3 + (ko*se - qo3)*dt/ks;
    gotoxy(i,24 - sr/2.0 - 0.5); putchar('r');
    gotoxy(i,24 - qd/2.0 - 0.5); putchar('d');
    gotoxy(i,24 - qo/2.0 - 0.5); putchar('o');
    gotoxy(i,24 - sp/2.0 - 0.5); putchar('p'); }
(void) getch();
}

```

The delay is put into the output function by the following lines:

```

qo = qo1;
qo1 = qo2;
qo2 = qo3;
qo3 = qo3 + (ko*se - qo3)*dt/ks;

```

The three dummy variables qo1 through qo3 provide a delay of 3 iterations. On the current iteration, the value of qo3 is computed just as it was in Part II. But this value is passed, on the next iteration, to qo2, then to qo1, and only on the third iteration to qo, the actual output quantity. So the output is always based on the error signal that existed three iterations, or $3*dt$ seconds, ago.

This is called a "transport lag." The variations in qo follow the variations in qo3, but three iterations later. The effect is much like shouting commands from one end-zone of a

football field to a marching band in the other end zone. The marchers follow all the commands in properly-spaced sequence, but with a delay about about 0.3 seconds.

It may not be obvious, but in the program of Part II there was already a transport lag of dt , one iteration. The effect of a disturbance of the input quantity during one iteration did not make its way around the loop to affect the input quantity until one iteration later. We saw that it was possible, with a delay of 0.1 second, to adjust the system for tight control despite this lag.

Now the situation is slightly different. Effects still propagate around the loop in one iteration, but the transport lag inserts a delay of 3 iterations in the output function alone. In ALL the functions, including the output function, there is a new value of input and output calculated on every iteration; the output quantity and input quantity can change during the transport lag period. Previously nothing at all could happen during the lag of one iteration. Now, with a lag of three iterations, the error signal can change on every iteration even though the effect on the output quantity is delayed by 3 iterations.

Compile and run the program. The appearance is much the same as in Part II, but now a "reaction time" is visible. After the initial rise in reference signal, the output quantity, represented as an 'o', does not start to rise until the fourth iteration, and the perceptual signal starts to rise one iteration later. The perceptual signal, a 'p', then rises smoothly until it matches the reference signal. When the disturbance jumps to 9.0 units, the perceptual signal is affected immediately. But the output signal, which opposes the disturbance, begins to change on the fourth iteration after the start of the disturbance. Then the perceptual signal begins to return toward the reference signal (some apparent lags are simply due to the coarseness of the resolution of the plot).

We have paid one penalty for this transport lag. The slowing factor, which had an optimum value of 10.1 in Part II, now must have a value of at least 70 to avoid overshoots. This added slowing is required in order to compensate for the transport lag and keep the control system stable. If you reduce k_s to speed up the transitions, the perceptual signal will begin to overshoot and undershoot. Try reducing the slowing factor k_s in steps of 10. The instability will get worse and worse until a runaway oscillation starts and control is totally lost.

Input delay

Now we restore the output lag to zero (or one iteration) and use the same mechanism to put a lag into the input function.

```
/* primer4.c */

#include "stdio.h"
#include "conio.h"

void main()
{
float sp = 0.0, /* initialize signals and quantities */
    sr = 20.0,
    se = 0.0,
    qo = 0.0,
    qi = 0.0,
    qd = 0.0;
float kd = 1.0, /* set constants */
```



```

        ki = 1.0,
        ko = 100.0,
        kf = 1.0,
        ks = 70.0,
        dt = 0.1;
float sp3 = 0.0, /* initialize variables for perceptual delay */
      sp2 = 0.0,
      sp1 = 0.0;
int i;
clrscr(); /* alternative: for(i=0;i<25;++i) putchar(0x0d); */
for(i=1;i<=80;++i)
{
    if(i > 40) qd = 9.0; else qd = 0.0;
    qi = kf*qo + kd*qd;
    sp = sp1;
    sp1 = sp2;
    sp2 = sp3;
    sp3 = ki*qi;
    se = sr - sp;
    qo = qo + (ko*se - qo)*dt/ks;
    gotoxy(i,24 - sr/2.0 - 0.5); putchar('r');
    gotoxy(i,24 - qd/2.0 - 0.5); putchar('d');
    gotoxy(i,24 - qo/2.0 - 0.5); putchar('o');
    gotoxy(i,24 - sp/2.0 - 0.5); putchar('p');
}
(void) getch();
}

```

Compile and run this program.

At the start of the run, when the reference signal is set to 20.0, the output immediately begins to rise. Examining the block diagram of the system (Part I), you can see that a change in the reference signal shows up immediately as a change in the error signal, which drives the output. This change in the output, however, is not reflected in the perceptual signal until the fifth iteration (it is probably affected on the fourth iteration, but not enough to show on the plot). When the disturbance switches to 9.0, the perceptual signal does not show the effect for three iterations even though the input quantity is immediately affected; then, when this effect does appear in the perceptual signal, the output begins to change at the same time.

So there is some difference when the delay is shifted to the input side. But control remains just as good as before, and with the same slowing factor of 70.

Sampled Control

Control systems can be designed so that they sample the state of the input at intervals instead of continuously. In the final program in this section, we use a timer variable *t* to count iterations. On every fourth iteration, the perceptual signal is set to the appropriate level on the basis of the input quantity. Between samples, the perceptual signal simply holds the last sampled value. All the other variables in the loop change on every iteration as usual. Here is the program:

```

/* primer5.c */

#include "stdio.h"

```

```

#include "conio.h"

void main()
{
float sp = 0.0, /* initialize signals and quantities */
    sr = 20.0,
    se = 0.0,
    qo = 0.0,
    qi = 0.0,
    qd = 0.0;
float kd = 1.0, /* set constants */
    ki = 1.0,
    ko = 100.0,
    kf = 1.0,
    ks = 70.0,
    dt = 0.1;
int i,t = 0;
clrscr(); /* alternative: for(i=0;i<25;+i) putchar(0x0d); */
for(i=1;i<=80;+i)
{
if(i > 40) qd = 9.0; else qd = 0.0;
if(t == 0)
{
qi = kf*qo + kd*qd;
t = 4; }
t = t - 1;
sp = ki*qi;
se = sr - sp;
qo = qo + (ko*se - qo)*dt/ks;
gotoxy(i,24); putchar('-');
gotoxy(i,24 - sr/2.0 - 0.5); putchar('r');
gotoxy(i,24 - qd/2.0 - 0.5); putchar('d');
gotoxy(i,24 - qo/2.0 - 0.5); putchar('o');
gotoxy(i,24 - sp/2.0 - 0.5); putchar('p');
}
(void) getch();
}

```

Once again, smooth control is achieved with a slowing factor of 70. Notice that while the output quantity changes smoothly, the perceptual signal changes in steps, one step on every fourth iteration. Obviously no disturbance can be resisted if it occurs during the hold period; the opposing output can only begin to change when the next sample occurs. The average delay would be half the sampling period. As with all control systems, disturbances can come and go too rapidly to oppose. But natural control systems are adapted to the environment that exists; most natural disturbances, therefore, can be opposed before their effects become important.

Best, Bill P.

Date: Sun Nov 15, 1992 4:03 pm PST

Subject: interfaces

[Avery Andrews 161128.1028]

(Hank Folson (921114))

I like the story about interfaces. It strikes me from reading Beth Preston's paper in interactionism that some netters have copies of, that she goes along quite nicely for about 15 pages, then basically collapses into an image of discrete events chasing each other around from outside to inside and back. Maybe what PCT does instead could be described as explaining how come you get invariants (at various levels of abstraction) maintained _at interfaces_.

(Dag Forssell (921114.0001))

I think I have a better sense of what you mean by 'engineering science' now. I don't think linguistics has any uniform approach to modelling - my personal preference is to get as close to the engineering type as possible. E.g. when I sat down to think about how you could organize getting yourself a beer, I fully expected to come up with an argument for action-planning, but it didn't work out that way at all. But in the area of grammar, I think that real modelling is premature. My parser, I would say, is not a model at all (it's an implementation of a competence description), while Penni's Salix is an embryonic model. But should we work on it by trying to go down towards neurology, or up towards integrating it into a fuller picture of what people are doing when they are talking, figuring out how you can have formulas and creativity, both obeying (more or less) grammatical generations, etc. It seems to me that 'linguists' (including Penni, who finds it amusing to be called one, hence the scare quotes) all think up is the way to go, and that we are getting unduly criticized for this. I re-recommend that anyone who wants to complain about our approach to this stuff ought to take on the exercise of modelling navigating your way somewhere on the basis of a text. As far as the linguistics end of it goes, it could be done on the basis of what I managed to teach myself about the subject in 12th grade in 1966-67 - this is not a boast about how clever I am, but an indication of how little you need to know. Mastering Shieber & Pereira's book on Prolog & Nat. Lang. Processing would be quite enough.

The stuff about the effects of engineering on the mind is interesting -- I certainly thought that learning about surface integrals, etc. had an effect on how I looked at things, though without getting as far into it as an engineer would.

Maybe I can say more about theories later, but other commitments are building up.

Bill Powers (921113.1000)

Gotta be quick, but I think I see a problem here. You seem to be assuming that first the words get connected with imagined perceptions, then the imagined perceptions with actual ones, whereupon there is a sort of 'negotiation' process whereby the imagined perceptions can be altered to a certain extent if they don't fit the actual ones.

The problem is that this alteration is constrained by the text: if I say that you get to the top of the hill and there's a large tree there, and you get to the top and see nothing but bare rock, with not even a stump or a place where one might have been pulled out, you really know that you've somehow got to the top of the wrong hill, or I've been lying to you.

So given that there is a system of constraints connecting texts with perceptions that might fit them, it's not clear that the imagined perceptions are doing any work. Why not just require that the actual perceptions fit the text directly? Or if not the spoken text, some lightly edited variant thereof, such as a portion of a KB?

What's nasty about this is that in effect the reference 'signals' (the text) and the perceptions (images) are things of different kinds, and we seem to be pretty clueless

about how to model the required kind of comparison process. But it has to be done anyway, so why not get rid of the middle-man? Where the imagined perceptions actually do some work, I think, is checking texts for consistency, which is a useful way to make sure you're resolving ambiguities, etc. correctly.

Avery.Andrews@anu.edu.au

Date: Mon Nov 16, 1992 9:57 am PST
Subject: Misc replies

[From Rick Marken (921116.900)]

Martin Taylor (921113 17:10) --

>Rick Marken wants a recap of the Vallacher and Wegner paper on "What do people
>think they are doing." Here are the relevant postings from Mark Olsen, Rick,
>and me.

Thanks.

The article seems to be about where people direct their consciousness after performing control activity.

Mark says:

> The relationship between attention and location of error in the hierarchy
>is evident here and they point this out.

Martin says:

>A gross simplification of their claim is that if the actions that allow high-
>level control are easy, then what subjects see themselves as doing is what
>we would call satisfying the high-level reference.

This is all very nice but the Vallacher/Wegner study is not really relevant to what I had in mind when coining PCT slogan # 1: "You can't tell what a person is doing by just looking at what they are doing". The slogan refers to the fact that an observer sees only observable aspects (sometimes relevant actions, sometimes irrelevant side effects) of the process by which a person controls their own perceptual experience (the behavior of these perceptions being what we mean by "doing" in PCT). It seems to me that Vallacher/Wegner assume that they already know what a person is doing and are interested in how these doings are described.

Ed Ford (921114:1115)

>Any time I deal with anyone, whether in private practice or elsewhere,
>What I was explaining in Section B was the process, the standards
>to be considered and incorporated when teaching others how to obey
>rules or helping them learn how to function effectively in an
>environment where rules exist.

I just don't get it, Ed. What does "teaching standards" have to do with a PCT based view of human nature? What I get from PCT is the idea that nothing could be less important -- the actual substance of a persons references for relationships, programs, categories,

principles, "standards", etc matters only in terms of how these satisfy higher level goals. The system should just be error free -- and this happens by having working (conflict free) control systems. Of course, such systems will be setting the "right" references for perceptions like your "standards", but they are right from the perspective of the control systems (they combine appropriately with prevailing circumstances to achieve the higher level goals. What is at any time a "right" setting for a particular standard from the point of view of the control system may very well appear to be a wrong setting from the point of view of someone who "knows the right standards". I know that the people you are dealing with have interfered seriously with other people in their efforts to achieve their goals. So, obviously, your goal is teach them to act without hurting others -- ie. "follow the rules". I think this is great -- but you should be clear that this focuses your treatment strategy on getting a person to act in ways that are better FOR YOU -- and, incidentally, for the person him/herself.

Avery Andrews (161128.1028)

>Gotta be quick, but I think I see a problem here. You seem to be
>assuming that first the words get connected with imagined perceptions, then
>the imagined perceptions with actual ones, whereupon there is a
>sort of `negotiation' process whereby the imagined perceptions can
>be altered to a certain extent if they don't fit the actual ones.

When we hear speech (in a language we understand) we (I anyway) hear the meanings (among other things) which are imagined perceptions -- not always visual. The "check out" doesn't always happen -- what ususally happens (I think) is an attempt to perceive that the person you are talking to is understanding (imagining) what you are imagining (and trying to communicate with words).

Best regards Rick

Date: Mon Nov 16, 1992 11:37 am PST
Subject: Re: PCT solgan #1; standards

[Martin Taylor 921116 13:00] (Rick Marken 921116.900)

> the Vallacher/Wegner study is not really relevant
>to what I had in mind when coining PCT slogan # 1: "You can't tell what a
>person is doing by just looking at what they are doing". The slogan refers
>to the fact that an observer sees only observable aspects (sometimes
>relavant actions, sometimes irrelevant side effects) of the process by
>which a person controls their own perceptual experience (the behavior of
>these perceptions being what we mean by "doing" in PCT). It seems to me
>that Vallacher/Wegner assume that they already know what a person is doing
>and are interested in how these doings are described.

I never had a problem with the intent of your slogan, but, as you nicely put it in this paragraph, you can't tell what your slogan is doing just by looking at what it is saying. V/W may or may not "know" what their subjects are doing, but the subjects say that they know, and what they say can be quite different without changing the externally visible actions. That's why I said V/W would agree with you, even without understanding the PCT implication. I don't propose any better wording. To understand what you mean is to understand PCT, so I think a succinct catchy slogan cannot really convey its intention except to one who already understands.

> What I get from PCT is the idea that
>nothing could be less important -- the actual substance of a persons
>references for relationships, programs, categories, principles, "standards",
>etc matters only in terms of how these satisfy higher level goals. The
>system should just be error free -- and this happens by having working
>(conflict free) control systems. Of course, such systems will be setting
>the "right" references for perceptions like your "standards", but they
>are right from the perspective of the control systems (they combine
>appropriately with prevailing circumstances to achieve the higher level
>goals. What is at any time a "right" setting for a particular standard
>from the point of view of the control system may very well appear to be
>a wrong setting from the point of view of someone who "knows the right
>standards".

I suspect this is an angel-free tread zone...

It seems to me that "standards" have evolved in much the same way as linguistic conventions. To some extent, it doesn't matter what they are, so long as people agree on what they are. They provide a simple way that conflict can be held to a low value, rather than having every interacting pair of people carry out a dance to discover a system of mutual accommodation. The concept of "evolutionary stable system" comes to mind. Some sets of potential "standards" cannot survive in any real society, because to use them would be to destroy the society--"kill whoever you see on sight" would be one such. Other sets are stable because deviation from them tends to set up conflict in the person who deviates. Such sets tend to be acceptable to the majority of the community, and deviation from them causes conflict that is enhanced by the actions of other community members that object to the deviant actions (they can't see the behaviour, by PCT slogan #1). So, conflict in each community member is reduced by conformity to community standards.

Problems arise when communities with different standards intermingle and their members interact. Then, adherence to either set of standards induces conflict, and there has to be a new dance to discover ways of reducing the conflict in both parties together. New "standards" must evolve that either destroy the old ones, or that allow the use of some of the old ones with the addition of standards that apply to the interaction of people from different communities. We used to call that diplomacy. Now we call it racism.

Martin

Date: Mon Nov 16, 1992 3:05 pm PST
Subject: From Mary Powers

[from Mary Powers (921116)]

I am resending Hank's message without the >'s because I don't know how to put them in. I have added or amended comments below each of his statements, beginning with a *** and indented.

[From Hank Folson (921115)]

Being interested in PCT applications, I came up with a list of 7 questions which are intended to be an application of PCT principles. The goal is to efficiently evaluate other theories, in case there is a better one than PCT. After all, there has to be something better than a theory that says: 1. You want to control everybody because you

are an independent control system, and 2. You can't control anybody because they are all independent control systems. The world would be in a terrible mess if it ran on that basis.

***Indeed it would. But control systems don't want to control everything. They want to resist those disturbances which create an error. From infancy human control systems have their errors reduced by the actions of other control systems. They learn give and take, and cooperation, which I believe is the dominant mode of interaction, not conflict. Humans are social, and social rules are developed and maintained in order to avoid conflict. They break down when following the rules leads to big errors in individuals (as when the rules are stacked against them too long and too obviously).

Penni answered them about Interactionism, and here is my answering of them about PCT, with the reasons for the questions in []. Bill Powers can correct my sloppy scholarship.

>1. Will you give me a couple of sentences long grossly oversimplified outline of the theory? See Mary Powers(921111) or Dag Forssell(921112) [This just gives you a rough idea about what the theory covers. It may allow you to see that there is no relationship to your interests.]

***my (921111) was my best shot, for now.

>2. Who developed the theory? About when? William T. Powers, Clarke and about 30-40 years ago. [This gives you an idea about the background of the theorists and a time frame so you have some idea of how mature or complete the theory is.]

***Hank dropped a stitch. Powers (Bill), Clark (Bob), and McFarland (Bob) in the 50's. Powers solo since 1960. Two physicists and a psychologist.

>3. What did (s)he believe in before (s)he came up with the theory? Bill Powers was studying traditional psychology. [This is straight PCT. The question is looking for the error signals that drove the new theory.]

***Bill had taken some undergrad psych. Unimpressed. Also involved in non-traditional psych which he left when it degenerated into a cult and a scam. He was working as a medical physicist but the interest in psych was long-standing.

>4. What did his/her old belief lack that the theory offered? Traditional psychology is based on a stimulus-response model, which Bill felt did not fully explain behavior. PCT offered an explanation of how living organisms function. [Again straight PCT. How did the new theory reduce the error signals in the creators?]

***Their goal was to explain psychology in a way that met their scientific standards as physicists.

>5. Where did the basic idea of the theory come from? Bill had been involved as a technician in WW II with control systems for naval guns. He made a connection between mechanical control systems and organically based ones. [New ideas are the result of reorganization of our experiences, per PCT. Evaluating the meaning of this answer can be very difficult. What does aiming guns at airplanes have to do with how living organisms work?]

***The basic idea came from Cybernetics. Bill was trained to repair control systems (radar and sonar equipment, not guns) in WW II. Then went to college for his BS in physics. Then read Norbert Wiener's book. Eagerly anticipating further developments in cybernetics, he began working on his own, drawing on his Navy background, etc., while waiting to see what the big scientific guns came up with. He's still waiting. The idea of control as a fundamental and pervasive property of living systems, as a key to understanding their organization and function, has never been seriously considered by the eminent and authoritative Great Minds in the life, social and behavioral sciences.

>6. How well has the theory worked for those who use it?

Well, it has given Bill Powers an interesting avocation... Ed Ford and others are successfully applying it to counseling and education. Franz Plooiij has done impressive work in biology. And so on. [This question gets to the applications of the theory. The proof is in the pudding.]

***Applications, schmapplications, this is basic research. The only model that explains how consistent ends are achieved by varying means in a changing environment. The uses to be made of it are largely unknown - the problem now is to intrigue enough people to learn it and try it.

>7. Has any independent research backed up the theory's claims? My two favorites are: Studies of E. coli bacteria by a non PCTer showed that it functions quite well using a very simple control system. This is an evolutionary stepping stone for the more complex control systems in humans. Plooiij studied the nipple rooting instinct in monkeys from a PCT point of view. (In traditional psychology for "instinct" read "magic"). He found that the lactating monkey mommy has a temperature gradient across her breast increasing towards the nipple. The baby monkey has temperature sensors in its lips. A simple control system leads the baby monkey to move its mouth towards higher temperatures. When it bites down on the nipple, other control systems take over. [This finds examples of real world proof or application of the theory for your evaluation.]

***The people who studied bacterial chemotaxis (Koshland, et al) are unaware of the PCT interpretation of their data. There's good data all over the place - nothing wrong with observations. It's hard to get at, though - often all that is published are the results of statistical massaging while the good stuff (from the PCT point of view) is buried in a file somewhere.

Note that if these questions were asked about PCT 30 years ago, the answers would not have been as complete. The questioner would have to wait for years for confirmation, or

help to study/prove the theory, or decide for or against the theory with incomplete information.

***Thirty years ago Bill believed he had applied some interesting principles of control theory to standard psychological data. Only after many, many years of being ignored, misinterpreted, and actively resisted, did he come to realize that control theory is a revolutionary paradigm in the life sciences, a systems concept that invalidates fundamental beliefs and commitments, and as such has been largely unacceptable.

Theoretically, Hank Folson

***Mary P.

* * *

On another topic:

I saw an article in the NY Times magazine for 10-18-92 on Antonio and Hanna Damasio - talking about convergence zones in the brain where lower level perceptions are constructed into percepts like specific faces. All very interesting, but what caught my eye were the following paragraphs:

Scientists once believed that signals flow through the brain in one direction only, from low levels to high. But in the past decade, several researchers, notably Gary W. Van Hoesen and Kathleen S. Rockland, now both at Iowa, and David Van Essen, working at the California Institute of Technology, also established that high-level regions connect back to lower ones - providing feedback pathways by which convergence zones can reactivate the sensory cortices.

None of this proves yet that Damasio is right. Cautions Morton Mishkin, a memory researcher at the NIH, "we know feedback connections exist, but we don't know their function."

We might have a suspicion or two - any suggestions as to how to communicate with these guys - like suggesting they read Behavior: the control of perception - without stepping on their toes?

Mary P.

Date: Mon Nov 16, 1992 3:47 pm PST

Subject: standards

Rick Marken (921116.1300)] Martin Taylor (921116 13:00)

>I suspect this is an angel-free tread zone...

Aw. What the heck.

>It seems to me that "standards" have evolved in much the same way as >linguistic conventions.

We're talking about cultural evolution here, right? No genes for particular linguistic conventions of social standards in other words?

>Other sets are stable because deviation from them tends to set up conflict
>in the person who deviates. Such sets tend to be acceptable to the majority
>of the community, and deviation from them causes conflict that is enhanced
>by the actions of other community members that object to the deviant actions
>(they can't see the behaviour, by PCT slogan #1). So, conflict in each
>community member is reduced by conformity to community standards.

It sounds like you are saying that there are certain standards which are adopted because not doing so creates internal conflict. I interpret this to mean that only by setting their reference for certain standards perceptions at a certain level can people avoid internal conflict. This is certainly true on an individual basis -- a person who wants to perceive themselves as socially cooperative would be creating a big conflict for themselves if, for some reason, the reference for a perception with a socially accepted reference (like, wearing clothes in public) were changed to a different value. By I don't believe that there are ANY standards perceptions which, if controlled at a particular reference level, would be intrinsically internal conflict producing; conflict depends on what other perceptions a person is controlling and at what level. I think Ed believes that there are certain intrinsically intrapersonal conflict producing standards perception reference settings. Do you believe this too, Martin? If so, could you give me an example of such a reference setting for such a standard perception?

I do believe that there are settings for references for standards perceptions that produce interpersonal conflict -- there are LOTS of them. You (Martin) gave one example -- setting the reference for the perception of what people to kill to "whoever". Such conflicts occur because carrying out the purpose tends to produce disturbances to INTRINSIC variables in the other person; there is a biological basis to much (but not all) interpersonal conflict.

Regards Rick

Date: Mon Nov 16, 1992 4:14 pm PST
Subject: Re: pasking levels

[Martin Taylor 921116 16:45] (Avery Andrews 921114.1250)

>Surely people rarely in real life do anything half as hard as
>controlling perceptions at the Pasking levels. Mostly people control
>for the other person looking as if they were paying attention &
>understanding. When this does not seem to be happening, they do
>something to attract attention (grad students who I'm familiar
>with might say 'hey, did you get that', or something), but there
>are very strong social constraints on serious checking of
>comprehension.

What happens in any particular case would depend on the difference between reference levels and perceptual signals. If the talker perceives (believes) that if the listener hears, then the listener will understand, all the listener need do to satisfy the talker's reference for perceiving the listener to understand is to take an attentive attitude.

In the current version of Layered Protocol theory, there are three propositions relating to the transmission of any message at any level of abstraction:

- P1) The recipient has made an interpretation of the message
- P2) If P1, the recipient has made a correct interpretation of the message
- P3) It is not worth continuing to try to transmit this message.

P1 and P2 relate reasonably closely to perceptual signals, P3 to an output gain control and an error signal (P3 doesn't have any very clear PCT analogue and may have to be rethought). It is assumed by default that the Originator had a reference signal that P1 and P2 become true, and an original perception that they were not true. If P1 and P2 do become true, there is no remaining error, and therefore P3 becomes true. But P3 can also become true if the reference for P1 and/or P2 changes, which can happen if the Originator perceives the Recipient to have correctly interpreted the higher-level message to which this one is contributing. P3 can also become true if the parties get into difficulties, so that there is a continuing error--this is a "giving up" condition, about which we had some discussion a month or two ago.

So now we can look at Pask's levels and Avery's comments. We assume that what is wanted is an understanding of the message, which means that Pask's third level of agreement is required. But that doesn't mean that an outside observer will detect actions corresponding to Pask's feedback. It depends on (a) the gain or insistence of the Originator's highest level SCS (Structured Control System, or Protocol Node in LP terms), and (b) whether at the highest level the Originator believes P2 initially to be true. If (b) is true, then there will be no feedback at that level. All that may be needed is a low-level determination of whether P1 is true--the listener heard the words, and for that, all that is necessary is, as Avery says: "Mostly people control for the other person looking as if they were paying attention & understanding." That is what LP theory (or equivalently PCT), combined with Pask, says.

I don't think there are "strong social constraints on serious checking of comprehension" when lack of comprehension would probably lead to serious perceptual error in the Originator of the message. If I tell you "We must get out quickly, because there's a bomb about to go off" I don't expect you to sit there looking attentive. I expect you to move. That's third level Paskian feedback, and if I don't get it, I'll continue to try to get it by whatever means I can. I'm controlling for the perception of your moving, which I would take to be a consequence of your understanding, given my assumptions about a few of the percepts you are controlling for and their reference levels.

Martin

Date: Mon Nov 16, 1992 4:37 pm PST
Subject: Modeling and linguistics

[From Bill Powers (921116.1400)] Avery Andrews (161128.1028) --

>Why not just require that the actual perceptions fit the text directly?

You can't compare text with a visual image. Text is perceived as a string of letters or phonemes. A tree on top of a hill looks like a tree on top of a hill, not like a string of letters or phonemes.

To make the comparison, you must either continuously translate what you're looking at into text and compare it with the original text, or translate the text into an image of a

suitable kind and continually compare the image with what you're looking at. You have to have a reference signal of the same kind as the perceptual signal it's to be compared with -- that is, it has to be the right sort of perception.

The least ambiguous way to do this, it seems to me, is to translate the text into an image, which you can then compare with the actual perceived images. When I hear "large tree at the top of a hill next to the path" I construct something like an imagined top of a hill with an imagined big tree next to the path. As I travel along I keep comparing what I see against the reference image. I'm on the top of a hill, so that error is close enough to zero. But the biggest tree I'm looking at is sort of middle-sized, and it's farther from the path than my image says it should be -- this probably isn't the top of the right hill. Keep going.

After you've given me a long set of directions, I doubt that I could repeat the text back verbatim. But I could paraphrase, because as you gave me the directions I was building up a mental image. If you asked me to repeat back the directions, I would simply imagine the mental images in the order I built them, and describe them. I might use some of the same words, but I wouldn't have to. I might say "Got it -- big tree, on the right, hanging over the path, up at the top of a hill," whereas what you actually said was "Look for place on the path at the top of a hill where there's a large tree nearby." You might want to tell me the tree's on the left, and that it doesn't overhang the path, and thus tweak my mental image. But what I will carry along in my head isn't the exact words, but the images they led to.

If an ambiguity comes up, I might replay the words. Wait a minute, this is the top of a hill and there's a big tree next to the path, but it's on this side instead of that side. What did he say? "A large tree nearby." Oh, he didn't say which side. So if I revise my mental image, it will match what I'm seeing. I'm there.

RE: the engineering approach.

I would prefer to call it the modeling approach, but the modeling I do is like what engineers do so maybe it's the same thing.

The main thing about a model is that it describes underlying mechanisms, not surface appearances. The model is something such that when it operates according to the way you defined it, its behavior creates the phenomenon you're trying to model. This is why, in a control-system model, you won't find any "controller." Control is a phenomenon that arises when you hook up an input function, a comparator, an output function, and an environment in a particular way.

What modeling is NOT is simply describing how things seem to work. Harris's operator grammar is not a model. It's a description of a phenomenon, in which some words (operators) appear always to be used with other words (arguments). That is what we have to explain with a model, but it's not a model in itself. It's a proposed description of a phenomenon. A model would provide the underlying mechanisms that lead to this result. And in that model we would not have "operators" or "arguments". Instead, we would have elements which, when they interact in the ways we propose, would create words that function as operators and others that function as arguments, and such that the way they are used fits the observations.

In a way, modeling as I see it is much like the kind of programming you probably do in implementing the grammatical or syntactic rules that you propose. When you're programming, you have in mind a certain phenomenon that you're trying to reproduce, and your problem is how to get a computer to do that. To get the computer to parse a

sentence, you can't just write the instruction "parse." You have to think of a series of operations which, when carried out, will produce something that you recognize as "parsing." None of those operations is parsing; parsing is an emergent phenomenon relative to the program steps. Nowhere in those steps will there be any operation called "parsing."

The difference in my kind of modeling is that I'm building a machine that is supposed to do what is required, not using one that's already defined. My machine is made of input functions, comparators, memory, and output functions, and its parts aren't all fully defined yet. When I try to construct an explanation of a phenomenon, however, I use the machine that I've been building, patching it up as necessary to get it to run. Those patches are to be treated with suspicion because they're ad-hoc, and one has to keep in mind that they require further work before they really become part of the machine. That's why I mention "unofficial" parts of the model now and then. But basically the process is done within the constraints of the conceptual machine components that are available.

The modeling process is then that of hooking up components of the kind that have already been defined to see if they can be made to behave in a way that matches the phenomenon to be explained. Their behavior isn't a matter of choice; that's a fixed property of the model. The only choices are in how we connect those parts. Once we commit to a connection scheme and define properties like perception and action to fit a particular case, the model will behave as it behaves. If the result doesn't match the phenomenon, we have to try again. We may have to find a better hookup; we may have to add new components and modify old ones.

When you use a computer to implement a rule or process, you're using a machine whose parts are already defined: adders, subtractors, index registers, logical comparisons, storage, retrieval, and so forth: the basic instruction set. You probably use a higher-level language, so each "elementary" operation is really a collection of machine-language instructions. But at whatever level you're programming, this model is still a digital computer. Your basic proposal is that the phenomenon is created by the operation of a digital computer. Or, at a higher level, you're saying that the underlying mechanisms consists of the basic operations you can perform in C, or Basic, or Pascal, or Common Lisp, or Prolog.

The conventional models of behavior that I have seen have been heavily influenced by the fact that the basic model is a digital computer or a language with fixed properties running on a digital computer. The influence of Lisp on AI is enormous; you could say that AI has been exploring the properties of an organism whose components are the basic operations of Lisp. Most conventional models have assumed that once a command is given, it will be carried out exactly -- because that is how computers work. The operations of these models have been conceived as tasks that are carried out stepwise, in sequence -- because that is how computer programs are written. It's assumed that one task can't begin before the previous one has ended, and that output always occurs after input. At best, you get time-sharing or subroutines. And that is because the underlying model is a digital computer or a programming language.

PCT is based on a different kind of machine having different components. It doesn't yet have all the components it needs, but that's just a matter of time. Each component, once defined, can be simulated on a computer, but the functions of the components are not dictated by what computers or programming languages can do. And there is no suggestion that the computer operations involved in such simulations are the same as the means by which the organism carries out the same operations.

To make a model of language, the first thing that has to be done is to get away from analyzing specific words, and to look for machinery that can handle words in general. There are certain basic operations on words that constitute the machine that produces language phenomena. These operations have to be independent of what the words are or what they mean, because they have to work with all words. There are certain basic phenomena that need modeling; these are the phenomena with which most linguistic models I have seen start, and which they seem to treat as explanatory. Given (father-of John Penni), what kind of underlying machinery can convert a perceived relationship into that or any related way of denoting the relationship in words? The underlying operations are clearly taking place in the linguist. But we need to get them into the model, if we're to have a true model. The "knowledge base" is the tip of the iceberg. What we need to know is what kind of machine can construct a knowledge base out of experiences?

Text appears to constrain experience. But that is just a surface appearance. To explain these appearances, we have to look at a deeper level, to operations that are of a more fundamental nature than word-manipulations. PCT is an attempt to construct that kind of model.

Best, Bill P.

Date: Mon Nov 16, 1992 4:37 pm PST
Subject: Formatting

[From Dag Forssell (921116.1415)]

Feedback is wonderful. Gary just posted the CSGINTRO back to me as a quality check. I am glad he did. Many of the lines were cut off with an extra hard return, with an annoying number of little orphans as a result. No important formatting lost this time, though, so my error signal is not nearly as large.

Gary's re-posting appears to allow 74 columns plus a hard return. I had used up to 77 plus hard return, leaving a margin of only two to the screen width of 80. Obviously, I shall re-format and re-post, unless this was an accident, caused by margin settings in Gary's word processor.

Before I do, I thought I would ask an open question of both Gary and those knowledgeable on the net if there is a standard number of characters that are desirable or maximum for the line length.

Bill, thanks for notes on file server.

Best, Dag

Date: Mon Nov 16, 1992 4:38 pm PST
Subject: Re: standards

[Martin Taylor 921116 17:30] (Rick Marken 921116.1300)

Angel-free tread zone ahead--->

We are indeed talking about cultural evolution in the setting of standards. No genes need apply (though some probably will).

>It sounds like you are saying that there are certain standards which are
>adopted because not doing so creates internal conflict.

Adopted by an individual, yes. But the point is a little more subtle than that. You are correct, but incomplete when you say:

> I interpret
>this to mean that only by setting their reference for certain standards
>perceptions at a certain level can people avoid internal conflict.

When there is conflict, there may be reorganization, and as Bill has often pointed out, that reorganization will tend to drive the conflicting systems into less conflict. If I do not conform to your standards, we both experience conflict if you care enough to try to make my actions conform (you can't see what I am "doing" but you can see my actions), and if your efforts make me unable to satisfy some references. So, point 1, it is not just me that experiences conflict and may reorganize. You may, too.

If a community has developed/evolved a set of standards that results in low levels of conflict when everyone adopts those standards for their actions (again, not for what they are "doing"), the standards will be rather stable. They work, because whatever people are "doing", their actions permit them to control their percepts adequately. That's what is meant by low levels of conflict. If the "standards" don't have this effect, and people find that they experience high levels of error when acting according to the standards, some people will reorganize one way, some another (it's random, after all), and the standards will disintegrate, perhaps to re-form as a new set of standards that provide lower overall error rates. Sets of standards that lead to sustained high error levels in many people are not stable. So point 2 is that if many people adhere to standards, it is because those standards do not conflict with the ideal of low intrinsic error.

> By I don't believe that there
>are ANY standards perceptions which, if controlled at a particular
>reference level, would be intrinsically internal conflict producing;
>conflict depends on what other perceptions a person is controlling and
>at what level.

I agree that there probably are no standards that we could call "intrinsic," but there are probably some reference levels that cannot be components of stable community standards. These will not be found in the standards of viable communities. But sets of standards probably fit together in clusters that are stable as a group that can be taken into or left out of a total system of standards. Different sets of precepts based on the teachings of long-lasting religions probably form such groups. I would imagine that the number of such sets that could be stable is unlimited, but the societies of the world may have found only a few tens of them.

If an individual lives in a community with stable standards, but does not use them to set the relevant reference levels, that individual will find conflict in many of his/her interactions with other members of the community, whereas the other members will find conflict only in interactions with the deviant. The deviant is more likely to reorganize than are the other members, and if the standard set is truly stable, this reorganization will continue until the deviant acts according to the standards and at the same time finds his/her reference levels generally attainable by non-deviant actions.

I suspect that most sets of social standards are not truly stable, and perhaps there are no possible sets of standards that lead always to zero error as a consequence of interactions. In a non-stable, or conflict retaining set of standards, all members of

the community are liable to reorganization, and the standards themselves will drift in a way directly comparable to linguistic drift and for the same reason. The result may be the breakoff of heretical groups, or a more or less unconscious shift of mores, or other shifts.

Serious problems arise, as I said earlier, when individuals who belong to different communities with incompatible standards have to interact. The incompatibility of the standards sets is defined by the existence of conflict when one individual uses one set and the other the other set. One or both must reorganize. When you have large numbers of individuals from each community meeting, then either one community will lose its standards (its "culture") to the other, or both will have to develop supplementary standards to deal with the interactions. That way lies stereotyping of members of "other" communities, but it may be a necessary way to handle the modern possibilities for world-wide interaction.

It's all based on the iterated interactions of individuals, and one-on-one reorganization based on the conflict that occurs.

Martin

Date: Mon Nov 16, 1992 6:26 pm PST
Subject: Re: pasking levels

[Avery Andrews 921117.1125]

(Martin Taylor 921116 16:45)

I'll have to mull this over, but a random remark:

>I don't think there are "strong social constraints on serious checking of
>comprehension" when lack of comprehension would probably lead to serious
>perceptual error in the Originator of the message.

I'm not sure, but I bet there are plenty of cases where nasty consequences follow from (some ?) people's reluctance to really make sure they've been understood, tho probably not in cases as clear cut as you bomb example.

Avery.Andrews@anu.edu.au

Date: Mon Nov 16, 1992 7:54 pm PST
Subject: Promoting PCT

[From Dag Forssell (921116.1940)]

I think this particular piece is in good shape now. The next step is to merge it with my demo tape. The interested observer may note the disappearance of the four kinds of theory. Good for discussion on this net, but unnecessary. Instead, I have developed a contrast with conventional training and consequent economic advantage. I have tried to be even more careful and specific in the claims, without backing off on my convictions. Engineers are not promoted nearly as much here. A cover letter or phone call is the place to suggest starting with the engineers. I will appreciate any comments.

UNDERSTANDING A THEORY-BASED LEADERSHIP PROGRAM

Summary

This paper suggests that a program based on a valid theory of human action is more effective than conventional training, which has no clear and valid theory at its core. Conventional training attempts to give rules for behavior. But rules for behavior are not appropriate. Behavior must be totally flexible to achieve consistent ends under widely changing circumstances. Human behavior is obviously flexible. A revolutionary theory shows how and why by explaining control and self-direction in detail. The new theory gives effective guidance in a large variety of settings.

Content:	Page
Summary	2
The value of a good theory	3
Objective "REALITY" versus personal "reality"	3
Individual development	3
Existing training programs	4
Theory-based education and training	5
Speed, cost, effectiveness	5
Scientific revolutions	5
One example of a scientific revolution	6
Perceptual Control Theory	6
Who will pay attention?	7
Who can understand?	7
PCT 1 applied in values, methodology and role play	8
Some of what you will learn	9
Relevance to work life	9
Relevance to personal life	9

1 PCT is an acronym for Perceptual Control Theory.

BOLD: The value of a good theory

Bertrand Russell, the famous British philosopher, is reputed to have said:

There is nothing as practical as a good theory.

In the engineering and physical sciences, this is obvious. Engineers and physical scientists recognize that a good theory allows for the prediction of performance long before actual experiment or production. Good theories have allowed us to travel, communicate, understand and produce better than ever before.

The power and practicality of a good theory is awesome. It also provides clarity.

When it comes to the important area of human affairs, this is not quite so obvious. Many theories have been offered over the years, attempting to explain human action, but none have measured up. This is why many psychologists say that their theories and practices have nothing to do with each other.

Since the results have been poor, scientists have settled for descriptions and statistical correlations between described phenomena. Explanations which yield predictions with 99+% experimental confirmation have not been possible and are not expected. Levels of explanation and proof required for publication are orders of magnitude lower than in the engineering sciences.

Some say this is because humans are too complicated. Could it be that the concepts have been inadequate? As we shall see, the history of science recounts many phenomena that were considered mysterious and too complicated until an appropriate concept was offered.

A good theory makes real breakthroughs possible.

BOLD: Objective "REALITY" versus personal "reality"

Let us distinguish between the objective physical REALITY of the world, and any one person's subjective interpretation of it: personal reality.

REALITY: What actually exists (but can never be known).

reality: What a person sees, hears, touches and smells.

Personal, subjective reality is like a map of REALITY. The personal reality is all a person can know. I think of reality as a heads-up display in a fighter pilot's helmet, except it is internal to our own head. It is only a display. REALITY is beyond our grasp, but we spend a lifetime trying to understand it. As we do, the subjective map of reality is formed in the brain and constitutes a personal world.

To a person who looks at a seamless world in living color and stereophonic sound, this distinction may seem silly. Obviously, what we see, hear, touch and smell is REAL, isn't it? Careful thought tells you that it cannot be. Your senses may be well calibrated and yield a very good map of the physical world, but a moments reflection tells you that it can only be a map. All you know is an accumulation of interpretations of signals from nerve endings. Today we know that there are many things in existence which we as humans can not sense at all. Elephant-talk below 18 hertz, echo-location beyond 20,000 hertz, infrared light, X-rays; all are part of a long list of phenomena we cannot sense. Clearly, our sense of the world is just that, sensed perception, a display.

BOLD: Individual development

A person grows up and develops a personal understanding of how the world "works" from personal experience, supplemented by experiences told by others, seen on TV and read in books. Each person develops a unique framework of ideas, a reality which governs that person's life. These personal ideas will vary. The quality of the personal map reality will vary widely from person to person.

A personal map of high quality, a good reality, allows us to navigate successfully in that elusive REALITY the map portrays. A good theory is a good reality; a map of a specific aspect of REALITY, expressed as description, explanation and perhaps a set of clear rules for prediction that follow from the explanation. Many people can easily share the same good piece of reality. This is why the power of a good theory is awesome.

BOLD: Existing training programs

Companies spend millions of dollars on training.

To illustrate the variety, this list is taken from a recent CareerTrackR brochure:

Team Building: How To Motivate And Manage People™
Total Quality Management™
Getting Things Done™
Making Meetings Work™
The One Minute Manager Live!™
Selling Smart™
How to Delegate Work and Ensure It's Done Right™
Assertiveness Training™
Personal Power™
How To Give Exceptional Customer Service™
Negotiate Like The Pros™
How To Set And Achieve Goals™
Stress Management For Professionals™
Controlling Anger™
How To Deal With Difficult People™
Self-Esteem And Peak Performance™
Enhancing Intimacy™

These programs may each be very good, but none are based on a good theory of human action. They are based on "what seems to work" - anecdotal experiences by the authors, experiences from many sources and conclusions drawn by the authors. These programs of course each portray a reality, but in the absence of proven theories in the area of human affairs they cannot offer a universal framework of explanation. The focus and quality varies.

These multiple programs, offering multiple scenarios and stories, suggesting multiple prescriptions for "what/how to DO," make understanding and dealing with people far more complex than it needs to be.

In most cases, people have fun and like the training, but four or five months later, little has changed in the workplace. I believe a large reason for this is that the training is situational / anecdotal and focuses on "what/how to DO." Each participant is left to integrate the many disparate lessons of the training experience into the framework of their personal reality, such as it is.

People want "practical" seminars focusing on "techniques," "skills" and "tools." This is all they have ever been offered, because absent a good theory, that is all anyone has been able to deliver.

People ask: Show me what to DO (cause) so I will get results (effect). Our program shows clearly why this is fallacious, but that does not change the fact that this is what many people have come to expect.

Instructions on "what/how to DO" are valid only in a given set of circumstances. Typically a training scenario is carefully selected and told with drama and humor by a speaker. You are told what the circumstances were, what was done and what the results were. You imagine that the same thing will happen if you do the same thing. You feel euphoric as you imagine success.

A large part of the "what/how to DO" training does not really apply in individual cases because the world is full of varying conditions and changing disturbances. The lessons become irrelevant and are soon forgotten. Euphoria fades away.

If much training is ineffective, how can we suggest that Purposeful Leadership is of lasting value?

BOLD: Theory-based education and training

The major strength of the Purposeful Leadership program is that it explains and applies a new theory called Perceptual Control Theory (PCT). PCT recognizes and explains the phenomenon of control. PCT explains why and how people do what they do. PCT is based on neurology and clear, detailed and tested engineering concepts. PCT requires and offers scientific rigor with explanation and prediction.

1 See separate paper: Control; what it is, where it applies.

PCT is an engineering science of psychology that is easy to understand for anyone who pays attention to the step by step explanation and expects a clear explanation, prediction, test and 100% validity.

Once the phenomenon of control is observed and the detailed explanation understood, it will be seen that control is the fundamental organizing principle of life. Control is pervasive and can be seen operating at a microscopic level as well as at the macro level of human activity.

PCT explains a wide variety of phenomena of everyday experience.

BOLD: Speed, cost, effectiveness

Instead of using *multiple* programs, each one covering some aspect of human interaction, you can use *one* to understand yourself and others in some detail. Participants can decide that the theory is good, by testing it in their own life. Everyone can draw conclusions from one theory supplemented with specific information as required by specialized applications. The time and expense of training is dramatically reduced. This one education is effective even as jobs are rotated, because a good theory applies everywhere (if it is really good and valid).

If this new theory is so much better, why is it not widely known already?

The reason is that this theory is dramatically different from the prevailing descriptive science of psychology. It causes a scientific revolution.

BOLD: Scientific revolutions

Thomas S. Kuhn, a leading scientific philosopher, professor at MIT and author of: The Structure of Scientific Revolutions³, explains. From the book cover:

...Thomas S. Kuhn wastes little time on demolishing the logical empiricist view of science as an objective progression toward the truth. Instead, he erects from the ground up a structure in which science is seen to be heavily influenced by nonrational

procedures... .Science is not the steady, cumulative acquisition of knowledge that is portrayed in our textbooks. Rather it is a series of peaceful interludes punctuated by intellectually violent revolutions ...in each of which one conceptual world view is replaced by another....

Nicholas Wade, Science

3 University of Chicago Press, 1970. (In print 1992).

Thomas Kuhn introduced the term "paradigm" in this book and suggests that scientists schooled in a certain set of views adopt them as their personal "paradigms" (or reality) and then view the world through these paradigms - as if they were eyeglasses filtering information. The word paradigm means pattern. It is used to signify how we interpret a phenomenon; how we explain the world to ourselves.

In Kuhn's view, everyone is a scientist, and every world view (or reality) might be called a personal science. Everyone has some framework of ideas of how the world "works" and views the world through those personal paradigms.

When a new theory (normally an extension of existing theory) is offered to an engineering scientist, the course is clear. Study the theory based on existing paradigms (which have proven 99.999+% dependable and usually allow excellent comprehension). Predict something based on the new theory, perform the experiment and accept or reject the theory expecting 99.999+% performance.

When a radically new theory is offered to a scientist, understanding can be difficult. If the old and the new concepts are incompatible, it may literally be impossible to see the new theory through the eyes of the old paradigm.

BOLD: One example of a scientific revolution

To illustrate, let us go back to a well known episode of a revolution between a well established descriptive science of astronomy and a new, incompatible, engineering science of astronomy.

In the beginning of the 17th century, Ptolemy's model (AD 140) of an earth-centered universe was still fully accepted. This model was complete with crystal spheres to carry the heavenly bodies and ample room for heaven and hell beyond the eighth sphere, the firmament. Learned scholars who had studied and accepted this model and fitted it into their reality, were not open to a new and different explanation.

In 1543, on his deathbed, Copernicus published a model of a sun-centered universe. (Kepler published additions in 1609-18). In 1610 and more explicitly in 1632, Galileo supported Copernicus' theory, based on the first ever observations by telescope and the discovery of Jupiter's moons. He had trouble being published in the scientific journals of the day (church bulletins?), due to the requirement that his manuscript be reviewed by a group of his peers.

(It is of passing interest that on October 31, 1992, Pope John Paul II acknowledged the church's error in this matter).

We understand today that his ideas were incomprehensible to the (church) scholars of the day, due to what they already understood as their personal paradigm. Of course Galileo's

writings were irrelevant to their science as they understood it and not a constructive and welcome contribution to the state of the art.

So Galileo self-published. Despite the persecution he personally suffered as a result, astronomy started over as an engineering science and made rapid progress.

Isaac Newton's *Philosophiae Naturalis Principia Mathematica*, (1687) completed the revolution, but was resisted in the same way. According to Kuhn, it took fifty years for Newton's work to be fully accepted.

This well known revolution is not an isolated case. Kuhn describes many upheavals in several disciplines. The opportunity for a revolution arises when a current paradigm fails to solve problems and competing paradigms are offered to provide better explanations. A struggle of many decades typically takes place. Established scientists continue the development, application and defense of the existing paradigm as usual while outsiders and early converts champion a new one.

A reading of Kuhn's seminal book makes it clear that there is lots of room in our sciences of today for coming revolutions.

New information - on any subject - is always filtered by what you already understand. PCT itself explains why this is so. Where a person has existing convictions, conflicting information is either not comprehended or rejected.

A person without convictions on a certain subject is more open to new information.

This is why scientific revolutions typically originate from outside the scientific community which has accepted the present paradigm.

BOLD: Perceptual Control Theory

Perceptual Control Theory is a new engineering science of psychology, 35 years young from its inception. It offers description, explanation and prediction. Explanations which yield predictions with 99+% experimental confirmation are possible and are expected in time. Much development work remains to be done, of course. Tests to date show 95-98+% correlation in simple experiments (which we will duplicate in class), with the remaining 2-5% accounted for by expected imperfection of control: less than infinite loop gain, slow response and sloppy connections in the environment.

PCT offers an opportunity for a transition from a descriptive science of psychology where theory and application are worlds apart to an engineering science of psychology, where theory and application fit like hand in glove.

For contemporary journals of psychology to publish articles on PCT, the phenomenon of control must be understood in detail in a review by peers who have internalized an understanding and strong convictions based on a scientific method limited to the study of cause and effect. An understanding, endorsement and acceptance of PCT is obviously very difficult. PCT itself explains why.

On the other hand, PCT is immediately acceptable -intuitively obvious- to people without such understanding and convictions. It is easy to understand and immediately useful.

Without a good theory, every problem must be solved by trial and error. You have to learn a lot of rules for every conceivable circumstance.

With a good theory (in high school or engineering school), you learn the theory, then spend time on problem after problem to learn to recognize how the theory applies and get used to using it. What you remember is the theory, not individual solutions.

With Purposeful Leadership and PCT, you learn a good theory, then spend time with application after application to recognize how the theory applies and get used to think that way. What you remember is the theory, not individual "what/how to DO" solutions.

The power of a program based on a good theory is awesome.

There is nothing as practical as a good theory based program.

BOLD: Who will pay attention?

PCT explains clearly that the only thing that drives us are our concerns. For this reason, the program, theory and applications must be understood to have relevance to the concerns that are in the forefront of the mind of each participant.

Those participants who have concerns which are addressed by the things PCT promises to explain will want to pay attention. Those who are satisfied with their own understanding, skills and results have no reason to pay attention.

BOLD: Who can understand?

Because personal concerns determine what is of interest and relevant, the program is structured to provide immediate relevance and usefulness as much as possible, with maximum clarity:

Day 1: Introduction to theory, values and methodology
Effectiveness application

Day 2: Related concepts
Leadership applications

The first and second day require careful attention, nothing more.

Day 3: Biology / engineering details.
Structural details.
Computer demonstrations.

Again, attention and interest is the key. A willingness to think in terms of explanation, prediction and test is very helpful.

BOLD: PCT applied in values, methodology and role play:

The value of Purposeful Leadership is in the theory we teach. To make the theory come alive with relevance to the individual, it must be applied. As the Chinese proverb says:

I hear	- I forget.
I see	- I remember.
I do	- I understand.

PCT describes how an individual operates in all circumstances. (We illustrate both cooperation and conflict in active demonstrations). PCT makes no value judgements.

But most people adopt values which constrain their options in working with others.

We specifically define (using PCT) what we mean by:

- Supporting individual effectiveness
- Respect for the individual
- Balance: Appropriate level of direction and involvement
 - minimize conflict
 - maximize cooperation

Given an acceptance of these leadership values/objectives, we offer a comprehensive methodology which can serve as a guide in a great number of circumstances.

We illustrate the theory, using this methodology, in several scripted vignettes. These are read, broken down to show the use of the methodology and discussed among participants.

The variety of interpretations and applications in these role plays makes it clear that the methodology is NOT a rigid "what/how to DO" set of rules, but only a guide to help you systematically consider various aspects of what the theory teaches you. Situations and disturbances in real life vary so much that the action required to accomplish the same outcome is almost never the same. This is why it is fallacious to teach "what/how to DO."

Some participants will view the methodology as a comprehensive "what/how to DO" prescription despite our comments to the contrary. (After all that is what people have come to expect from seminars in a cause-effect world). There is no harm in this. Taken as that, the methodology is an excellent "what/how to DO" prescription.

Participants are then invited to plan and act out a sequence of similar situations, given a scenario which is relevant to them, taken from their world of daily work. This scenario is developed ahead of time by a representative of the group and the teacher.

Each role play for a participant provides practice for a work situation, the way a sales role play provides practice for the salesman's work. This makes it personally relevant. The scenario for role plays is tied to the work environment where the training is to be applied, and provides continuity between role plays, from one to the next. Additional details will come from the common experience and imagination of participants.

This detailed scenario will do triple or quadruple duty. It portrays (for each group) some typical problem or source of conflict between fictitious characters such as Joe, Dave and Bill which can:

- 1) be resolved between Joe and Dave in a manager / subordinate or peer to peer conflict role play, focusing on wants, then
- 2) be facilitated by Joe in a team building situation between Dave and Bill, focusing on mutual perceptions, then later
- 3) proudly reported as an accomplishment by the employee Dave or brought up as a concern by the manager Joe in a performance review and finally

- 4) serve as a concern Dave has, addressed by the salesman Joe. (Sales only if a sales application is desirable given a particular audience).

Role play 1 is on the agenda for Day 1, role plays 2-4 are on the agenda for Day 2.

With a well chosen role play scenario, participants will in effect begin to solve (if all parties to a problem are present) or plan how to solve those problems that are close to their hearts.

BOLD: Some of what you will learn:

- * How quality is one dimension of individual self-direction.
How and why all understanding is individually subjective.
How and why what you choose to want affects what you do.
How and why what you choose to perceive affects what you do.
- * How and why dissatisfaction is the only mother of invention.
How you can influence what another person wants without conflict.
How you can influence what another person perceives without conflict.
How the last two points can add up to cooperation,
and how this relates to the old maxim: The customer is always right.
- * Why you cannot tell what people do by watching what they are doing.
How and why all behavior controls perception.
- * What motivation is and why you cannot motivate another person.
How to develop and support individual effectiveness.
- * Why two people can look at the same facts and draw different conclusions.
How and why your memories affect what you do.
How habits are memories behaving.
A way to understand imagination.
A way to understand feelings. Do they drive you?
What it means to respect another person.
How and why stimulus-response is an illusion.
- * How and why a reward may not motivate but create resentment instead.
- * When helping a person conflicts with respect for the person.
- * How to communicate effectively with subordinates without conflict.
- * How to conduct mutually satisfying and productive performance reviews.
What is required to develop and maintain an effective team effort.
- * How to resolve conflict with another person if at all possible.

(* = bold to break up list and attract attention)

BOLD: Relevance to work life:

Conflict resolution: - in a superior / subordinate relationship.
- in a peer - to - peer relationship.
- in a subordinate / superior relationship.

Total Quality Management framework.

Leadership understanding.

Team development.

Performance reviews.

Vision/Mission statements.

Goal setting.

Non-manipulative selling.

Teaching effectiveness, initiative and responsibility.
Develop trust, mutual respect and high morale.

BOLD: Relevance to personal life:

Conflict resolution.
Develop loving relationships.
Encourage self-sufficient, capable family members.
Develop self confidence.

-End-

Best to all, Dag

Date: Tue Nov 17, 1992 3:40 am PST
From: Jackson
Subject: Hello Dag! ...Promoting PCT

Hello Dag...

I am trying to put the pieces together here in Phoenix, but I woke up from my grief so far behind in school, work, and home that I can hardly breathe until January. I'm totally inundated.

By the way, thank you for the thoughtful note regarding my personal loss; your comments were deeply appreciated. Anyways, consistent with the rest of my life, I'm WAY behind on the Net. But as I've mentioned before, I always follow your stuff.

My "sales job" at my plant (homework?)...don't get your hopes up (mine have dropped some); I can't detect an error signal in our execs. However, I'm still trying to grow the small, but enthusiastic, following I've found.

I'm glad you had some success with your 3-day program, but the criticism from the HR person doesn't surprise me (and it probably didn't you, either). We in education understand that we embody our own learning preferences in our teaching techniques, and you are certainly the consummate engineering type (I mean that as a compliment). Maybe I can help you slant your material from the manufacturing perspective...I would certainly welcome some insight on how to work with engineers as well.

Regarding the first draft of of "Promoting PCT", most of the the thoughts I had were pretty much covered with your second draft; however, I'm going to snail mail you some of my comments on that edition soon. Yet, I must say that the version I scanned today(#2) was a real improvement over something which I was really excited about in the first place. I'd call you with some of my input, but I usually don't get a chance to go over this stuff until the wee (very wee) hours of the night.

There are some other "education" things I'm struggling with from a PCT standpoint which I am going to try to put out (carefully) over the Net in the next few days or so; obviously, I'll be looking forward to whatever input you have as well.

I'll be in touch soon...wanted to let you know I'm still out here; also, I haven't forgotten that we need to settle up on the "Freedom from Stress" and sales program audio tapes.

Best Regards, Ray

Ray L. Jackson 3613 W. Saragosa St. Chandler, Az 85226 attmail.com!rljackson

Date: Tue Nov 17, 1992 5:50 am PST
Subject: for dag

Dag,

I always thought the quote you start your 'leadership paper' with was from the famous German Psychologist Kurt Lewin. Did kurt steal it from Bertrand Russel?

--Todd

Date: Tue Nov 17, 1992 8:22 am PST
Subject: Formatting

[Silvert 921117]

>[From Dag Forssell (921116.1415)]

>

>Gary's re-posting appears to allow 74 columns plus a hard return. I had
>used up to 77 plus hard return, leaving a margin of only two to the
>screen width of 80. Obviously, I shall re-format and re-post, unless this
>was an accident, caused by margin settings in Gary's word processor.

>

>Before I do, I thought I would ask an open question of both Gary and
>those knowledgeable on the net if there is a standard number of
>characters that are desirable or maximum for the line length.

The standard width (which is the default for fmt and other utilities) is 72 characters. This leaves room for additional characters like the ones used above to indicate material being replied to.

Lines that wrap are virtually unreadable, especially when preceeded by >. Try this for example:

>Before I do, I thought I would ask an open question of
both Gary and
>those knowledgeable on the net if there is a standard
number of
>characters that are desirable or maximum for the line
length.

--

Bill Silvert at the Bedford Institute of Oceanography
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2
InterNet Address: bill@biome.bio.ns.ca

Date: Tue Nov 17, 1992 9:31 am PST
Subject: Re: pasking levels

[Martin Taylor 921117 11:00] (Avery Andrews 921117.1125)

>I'll have to mull this over, but a random remark:

>

>>I don't think there are "strong social constraints on serious checking of
>>comprehension" when lack of comprehension would probably lead to serious
>>perceptual error in the Originator of the message.

>

>I'm not sure, but I bet there are plenty of cases where nasty consequences
>follow from (some ?) people's reluctance to really make sure they've been
>understood, tho probably not in cases as clear cut as you bomb example.

Nasty consequences can follow from any act, and it is not always possible to control the perceptions that they disturb. But if nasty consequences normally followed from some actions that were the result of error in soem ECS, the individual would probably have reorganized. "Nasty consequences" = high error -> reorganization.

If there are social constraints against ensuring understanding, either there would tend to be more error in other perceptions if people pressed for full understanding, or the social system (standard set) is unstable and evolving. Perhaps both. Social conventions presumably evolve toward actions that tend to result in minimum overall error in the individuals who use them in interactions.

Martin

Date: Tue Nov 17, 1992 9:52 am PST

Subject: Formatting -Reply

subject: formatting
(from Len Lansky 921117.1149))

I am not knowledgable about the system, but I do know that I messages, when downloaded to disks, do not fit properly. THus if the lines were, by convention, made shorter, I would like it.

I will check this message too and send a followup if it does not work. It is 64 units wide and fits a "message box" in our setup at the University of CIncinnati.

Len Lansky

Date: Tue Nov 17, 1992 9:52 am PST

Subject: Formatting -Reply

second reply to formatting
[from Len Lansky (921117.11570)]

I goofed last time by not using brackets and parentheses. Sorry about that.

My format fit my disk; Dag--yours did not. Often a word or two were dangling in next line. I suspect 64 or 65 is maximaum that will work with my setup. I would like to see such a norm set; it

would make my work much easier.

len lansky

Date: Tue Nov 17, 1992 10:23 am PST
Subject: foundations of linguistics

[From: Bruce Nevin (Tue 921117 12:28:49)]

(Bill Powers (921113.1000)) --

>-----
>We have quite a few different approaches to linguistics on this net.
>I've been wondering: is there anything about language on which ALL you
>linguists can agree? It seems to me that as in physics, there has to be
>a level of description that is so simple and low-level that we can at
>least agree on what phenomenon is occurring at that level. What are the
>absolutely non-controversial aspects of language that different people
>explain differently?
>-----

Bill, there is little unanimity and much contention among the various flags raised and defended in the field called linguistics. It seems clear to me at least that there can be none until the field has a firm scientific basis, which I believe it now lacks, and that in a PCT context we may establish such a basis.

Reinterpreting every proposal and every claim in terms of perceptions and control of perceptions is a much needed emetic. The patient has to swallow the medicine, however, and go away privately for a spell to reestablish equilibrium. This has to happen singly for each linguist, and on interim visits we tend to have a dissheveled appearance. Perhaps others who are committed to the process can share my sentiments when I say that I expect agreements in the outcome but probably not in its antecedents.

To continue the metaphor, I believe I have been on a rather spare health food diet for most of my career in linguistics, and consider myself fortunate in having relatively little in the way of extranea to lose. But de gustibus, and the fat lady hasn't done her thing yet.

Sorry about the imagery. Must be the weather. Bruce bn@bbn.com

Date: Tue Nov 17, 1992 1:06 pm PST
Subject: propositions

[From: Bruce Nevin (Tue 921117 12:46:29)] (Martin Taylor 921113 16:00) --

OK, so "propositions" are in this usage (you, Bill C.) not language entities but nonverbal perceptions. For me, "propositions" are sentences in a language, usually a formal "language" related to natural language (it is claimed) by some machinery of formal semantics. This is I think standard usage in philosophy and linguistics, but correct me if I'm wrong. I'll try to keep the two usages distinct.

As I recall, this started with my suggestion to Bill C. that the Hamming distance is not to be measured between sets of propositions (sentences in a formal "language") but between sets of perceptions. For computational purposes, these perceptions must be represented as propositions (=sentences). However, it is important to bear in mind that

they are not propositions, conceived typically as hanging out there in neutral-observer-space between the communicators. Instead, they are private perceptions within each communicator, about which they may severally formulate and utter propositions in some (natural or formal) language. They may come to agreement about the propositions (=sentences) at Pask's level 1 of repetition or level 2 of paraphrase. Agreement level 1 could be reached by a record/playback device. Level 2 agreement could be reached by a computer equipped with a natural language parser and capable of paraphrase. However, level 3 agreement cannot be reached by anything other than a control system controlling perceptions that correspond to those of the originator. Well, maybe some simulacrum could be reached with an elaborate enough knowledge base associated with the parser/paraphraser used for level 2, but I doubt this arrangement could pass the Turing test for long.

Pask's level 3 feedback gives to the Originator O (in your words)

>> an iron-clad guarantee that
>>R had received what O intended to send. Such an agreement seems to me
>>to be more than a simple agreement about a simple perception, and far
>>more than an agreement like "yes, we employed that string of words."

At level 3, R refers to his or her perceptions associated privately with O's message, and talks about those perceptions without either repeating or paraphrasing what O had said. Closing the loop, O finds that this description accords with what O might have said about O's private perceptions, and tells R so.

You said that the "propositions" are instances of the "primal message" in O's communication to R. But (in your words) "the Primal Message is a the reference state for the Originator's (O's) belief about the Recipient's (R's) belief." In other words, the "primal message" is a remembered or imagined perception. A belief is not a proposition (=sentence), though sentences may represent or express a belief.

An aside: a belief (it seems to me) is a perception that you assume (perceive) to be a true reflection of Boss Reality. But that is not sufficient. By that definition, when I look out my window and perceive that it is snowing I would have to call that perception a belief. Only a PCT aficionado would find no error there ;-). For most folks, a belief is an imagined perception that we perceive as veridical but unsupported by direct evidence. But most folks (that is, most of us most of the time) identify beliefs with statements of belief, just as they (we) confuse perceptions in general with descriptions of them. Must have something to do with how words evoke imagined perceptions and perceptions evoke words--however that happens. But I wander.

Bruce bn@bbn.com

Date: Tue Nov 17, 1992 2:32 pm PST
Subject: Re: standards, degrees of freedom

[From Rick Marken (921117.1000)] Martin Taylor (921116 17:30)

> If I do not conform to your standards, we both experience conflict

Not necessarily true. Here you are talking about interpersonal conflict; we "experience it" only in terms of the success (or lack thereof) of our efforts to control variables (in my case, perceiving you as conforming to my standards, in your case, perceiving no loss of control as a result of my efforts -- beating you, starving you, locking you up,

etc -- to get you to conform to my standards). If one person is a lot stronger than the other, they will "experience" no conflict at all in this conflict; they will just get the result they want. If both people are about equally strong, they will experience loss of control -- ie. error with respect to some variables they are trying to control. Of course, being people, they will also be able to perceive the cause of their lack of control -- the other person.

> But sets of standards probably fit together in clusters
>that are stable as a group that can be taken into or left out of a total
>system of standards. Different sets of precepts based on the teachings
>of long-lasting religions probably form such groups. I would imagine
>that the number of such sets that could be stable is unlimited, but the
>societies of the world may have found only a few tens of them.

I basically agree with this. My problem with this whole analysis is just the emphasis on "standards" setting as a basis for harmonious interactions in groups of control systems. I think this is almost certainly a crock. As humans we do happen to be able to perceive at the system level but that doesn't mean that controlling perceptions at this level is any more important than controlling perceptions at other levels. Herds of animals, for example, work together just fine without agreement on (or ability experience) standards, system concepts, principles, categories or whatever. Most everyday conflicts between people are usually over control of perceptions that are a lot lower level than "standards" -- and people work them out just fine.

I think organisms in groups "get along" when there are a sufficient number of perceptual degrees of freedom to be controlled -- and sufficient environmental degrees of freedom to allow all members of the group to control their perceptions. This means that the organisms must be able to perceive the environment in a way that allows simultaneous solution of the perceptual df problem in the constraints of the environmental df. Tom Bourbon's studies of two people controlling the relative distance between lines on a screen contributes more to our understanding of what makes it possible for multiple control system to "get along" than does all our blathering about standards setting. Standards are just one thing people have to be able to control -- no more or less important than controlling sensations, configurations, transitions, etc. When people can control their perceptions -- and when each individual in a group can control his or her own perceptions -- then there will be no interpersonal conflict -- no crime, no creeps, no assholes. This is an achievable goal -- but to get there we have to look in the right place; NOT at figuring what standards people should set, but at figuring out how to provide people -- ALL PEOPLE - with the degrees of freedom necessary to control their own perceptions. We already know how to do this, actually -- PCT just shows WHY this is important: The way to do it is

1) POPULATION CONTROL (to preserve the available df)

2) EDUCATION (to learn about the available df for controlling our own perceptions -- and how to control those perceptions more effectively).

People have tried to solve their problems by finding the right standards for centuries (from the beginning of recorded history) -- it not only doesn't work; it is the cause of most of our intractable problems (nationalism, religious wars, etc). I suggest that we approach the problem of interpersonal interaction from a PCT perspective; if people really are input control systems, then PCT should have some scientifically and practically useful things to say about how multiple control systems can get along without conflict. I think the answer is "degrees of freedom" -- not "standards".

Best regards Rick

Date: Tue Nov 17, 1992 3:51 pm PST
Subject: Re: "perception" or "belief"

[Martin Taylor 921117 16:40]
(Bruce Nevin Tue 921117 12:46:29)

>An aside: a belief (it seems to me) is a perception that you assume
>(perceive) to be a true reflection of Boss Reality. But that is not
>sufficient. By that definition, when I look out my window and perceive
>that it is snowing I would have to call that perception a belief. Only
>a PCT aficionado would find no error there ;-). For most folks, a
>belief is an imagined perception that we perceive as veridical but
>unsupported by direct evidence. But most folks (that is, most of us
>most of the time) identify beliefs with statements of belief, just as
>they (we) confuse perceptions in general with descriptions of them.

I have used "belief" to represent a structured perception. Since "classical" PCT takes the perceptual signal to be a single scalar value, and I identify "perception" with "perceptual signal" at least some of the time, I need another word to describe a complex of many simultaneous perceptual signals at a common level. "Belief" seems to me to be appropriate. A perceptual signal represents all you can know about the present state of a real or imagined CEV. A belief represents all you can know about a more complex set of states in the real or imagined world. And yes, when you look out of the window and believe you see snow, it is a belief that it is snowing, even though the flakes may be ash from a forest fire. I'm sure I'm not alone in finding that what I "directly" perceive is sometimes shown to be false after more observation.

If you have a better word than "belief" to represent a complex of perceptions, then I might well use it. But for now, it seems to relate to the everyday meaning of the word about as closely as the PCT meaning of "perception" corresponds to its everyday meaning.

As for the word "proposition", I'm not happy with appropriating it from its mathematical origin, but again, I have no better. Propose one.

Martin

Date: Tue Nov 17, 1992 3:56 pm PST
Subject: Formatting -Reply

correction to previous messages

[Len Lansky (921117.1619)]

Thanks to someone who wrote me directly about the issue, I double checked. I can cope with the length as it now exists IF I PRINT FROM DOS AND IF I READ IN DOS ON THE SCREEN. However, when I made hard copies from, or read the screen in wordperfect after downloading from bitnet, I had many "orphans" etc.

I wonder if others may have a similar problem with formatting.

Len Lansky

Date: Tue Nov 17, 1992 4:07 pm PST
Subject: Language

[From Bill Powers (921117.1430)] Wolfgang Zocher (921116) --

I tried to reply to you but there's something wrong with the address I have. In brief, your program sounds wonderful and I would like a copy of the source code (email should be OK). Let's work together on it for a while; it sounds as if it should be ideal for the Primer series.

Martin Taylor, Bruce Nevin (921117) --

Pask's level 3 seems to require the same finite universe of propositions that the Hamming-distance idea needs. You'll notice that the level 3 feedback was a formal statement about the properties of a circle. This would not work if the figure that was being described was a squashed irregular oval (in other words, some natural shape not clearly and formally distinguishable from others). There is no "iron-clad" guarantee that R has received what O intended to send unless the number of alternative guesses is small, and each guess is separated cleanly and by a large "distance" from all other possible intended transmissions. I don't think that is the general case in communication.

Bruce Nevin (921117) --

RE: ideas held in common among linguists.

You sound to me like a person who has done a lot of reorganizing about PCT. I think you understand what I mean when I say we have to start from scratch. I was talking about psychology and you're talking about linguistics, but it's the same phenomenon.

If linguists don't have any shared reality, then perhaps that leaves us free to invent one.

Best, Bill P.

Date: Tue Nov 17, 1992 4:13 pm PST
Subject: Re: standards, degrees of freedom

[Martin Taylor 921117 17:00] (Rick Marken 921117.1000)

> My problem with this whole analysis is just
>the emphasis on "standards" setting as a basis for harmonious interactions
>in groups of control systems. I think this is almost certainly a crock.
>As humans we do happen to be able to perceive at the system level but
>that doesn't mean that controlling perceptions at this level is any more
>important than controlling perceptions at other levels.

I'm a bit confused. I knew this was foolhardy territory to get into, but I can't see what the discussion of interaction procedures that evolve into conventions known as standards has to do with the system level. You seem to be implying throughout your response that the only place where standards exist is at the system level. I intend the term to apply at all levels relevant to interactions among people, and I think it applies probably more to actions than to behaviour (using the PCT distinction that behaviour is the control of one's own perceptions, whereas actions are not). Standards include greeting patterns, dress codes, thank-you notes for gifts, and all sorts of things for which the external appearance is what matters.

If I can act according to the standards of my community, and nevertheless control my perceptions with little error, I won't reorganize much, and will continue to act according to the standards. If I don't act according to the standards, and nevertheless am able to control my percepts (like the GodFather example you give), I won't reorganize. But in most cases, if the standards matter to many of the people with whom I interact, I will find that not acting according to the standards may impede my control (I may not get the job because I didn't wear a suit to the interview; I may not get a gift from Aunt Mabel because I didn't send a note thanking her for the last one), and I am likely to reorganize. When my reorganization leads me to act in such a way that I maintain control of my percepts, I will no longer reorganize.

Real community standards are those that tend to induce reorganization in people who don't act according to them. As I said before, their stability is determined by the degree of conflict occasioned on average in people who abide by them, because error will lead to reorganization and if there is a set of standards not very different from the current set but that tend to lead to less error, then the community standards will drift in the direction of that set. The word "community" is diffuse here. It is clearly weighted by the probability of interacting with any particular person, so for most people I suspect the standards one develops will be closest to those of the parents and older siblings, at least when interacting with them, though other sets of standards may be developed for interaction with others (such as the local gang).

> Standards

>are just one thing people have to be able to control -- no more or
>less important than controlling sensations, configurations, transitions,
>etc.

As you can see, I don't think standards are anything that people control (or other herd animals, for that matter). They are the products of reorganization, not percepts. They are the ways that percepts can be controlled when other people are involved in the actions that together form the controlling behaviour. All the same, I suppose that people can model desirable organizations, talk about them, and explicitly teach them to the young. But the problem here is a generic one with how you teach any behaviour deliberately. The "standard" that you can talk about is a model of a simulation, not the result of a structural reorganization.

> I suggest that we approach the problem of interpersonal
>interaction from a PCT perspective; if people really are input control
>systems, then PCT should have some scientifically and practically useful
things to say about how multiple control systems can get along without
>conflict. I think the answer is "degrees of freedom" -- not "standards".

Precisely what I am trying to do. "Standards" is the result, not the instigator. It is the manifestation of the dynamics of an uncontrolled interaction among control systems,

not a prescription for what should happen. It becomes a prescription by methods fully intelligible within classical PCT.

Does this seem even more at cross-purposes with what you are talking about, or does it make things clearer?

Martin

Date: Tue Nov 17, 1992 4:29 pm PST
Subject: Re: Language

[Martin Taylor 921117 17:30] (Bill Powers 921117.1430)

>Pask's level 3 seems to require the same finite universe of propositions
>that the Hamming-distance idea needs. You'll notice that the level 3
>feedback was a formal statement about the properties of a circle. This
>would not work if the figure that was being described was a squashed
>irregular oval (in other words, some natural shape not clearly and
>formally distinguishable from others). There is no "iron-clad" guarantee
>that R has received what O intended to send unless the number of
>alternative guesses is small, and each guess is separated cleanly and by
>a large "distance" from all other possible intended transmissions. I
>don't think that is the general case in communication.

I think you are only saying that it takes an infinite amount of information to describe a real number. Yes, it does, and it would take an infinite amount of communicated information to arrive at an agreement over an understanding on a proposition from a continuous universe.

I don't buy the "iron-clad guarantee" either. But that was what Pask claimed. I don't buy it because it is impossible to have error-free communication over any physical channel. The same holds for any perception. It's not unique to communication. All you can really talk about is the amount of information available in your perceptual value or equivalently your communicated message.

For most purposes of communication, we are working at the category level or above in the messages we want to communicate. So the universe is not only finite, but each "guess is [usually, MMT] separated cleanly and by a large "distance" from all other possible intended transmissions."

Martin

Date: Tue Nov 17, 1992 6:30 pm PST
Subject: Re: Modeling and linguistics

[Avery Andrews 921118.1311] (Bill Powers (921117.1430))

>If linguists don't have any shared reality, then perhaps that leaves us
>free to invent one.

You might, but I wouldn't recommend it.

Avery.Andrews@anu.edu.au

Date: Tue Nov 17, 1992 7:37 pm PST
Subject: Re: Modeling and linguistics

[Avery Andrews (921118.1430)] (Bill Powers (921116.1400))

>The least ambiguous way to do this, it seems to me, is to translate the
>text into an image, which you can then compare with the actual perceived
>images. When I hear "large tree at the top of a hill next to the path" I
>construct something like an imagined top of a hill with an imagined big
>tree next to the path.

This won't work, because images have the wrong kind of information in them. A classical example from philosophy: an image of a triangle has to either be isocoles or not isocoles, but the concept triangle is neither. It's got to be something that subsumes both. Or another example, an image of Ronald Reagan on a horse, unobscured, will have the horse white or not white, Reagan with or without a hat, etc., while if I say 'walk down the corridor until you get to the picture of RR on a horse', neither of these attributes matters.

It is indeed something of a mystery how anything can translate between images and the 'concepts' that we like to represent as propositions in logic, but that mystery has got to be solved anyway, & once it is, this solution could be used to build a concept-to-image comparator.

On the other hand, thinking about this stuff, it does strike me as quite plausible that one's interpretation of instructions could be quite strongly influenced by images formed at the time they were heard, so that what one did on the basis of a given set of instructions would be based on a confused mixture of what was remembered of the text, & of images & maybe other constructions formed on the basis of the text when it was heard. This looks like a possible subject for experimentation - I wouldn't be at all surprised if work had already been done on it.

Sorry to have to be brief, but busy day today.

Avery.Andrews@anu.edu.au

Date: Tue Nov 17, 1992 8:36 pm PST
Subject: Thought about Protocols

[gabriel 921171.21:31CST]

This is my version of Martin Taylor's protocols. I had not articulated them as protocols until Martin's post, so the original idea is his.

My baseline concept is a Kanerva database of truth values for a set of N Aristotelian propositions that together discriminate within an area of discourse. The truth values are bit strings N long, and can be thought of as being in a space of N dimensions, containing 2^N points each representing a possible object in the area of discourse. Typical values of N are in the range of 100-200, and the space is very sparsely populated.

The literature ref to Kanerva is "Sparse Distributed Memory; P. Kanerva, MIT Press 1988"
The book is not really hard, but it isn't easy either, and I still find some of the algorithms fading in and out of my understanding as I look at them.

Level zero is establishing agreement about the propositions that discriminate in the area of discourse. Most of them are already known to both parties and taken as given, and most of the the time they are all shared. It seems to me that one trouble we have in communicating about PCT is that our propositions defining or discriminating within the PCT area of discourse are not all shared by non PCT'ers. In fact I think we don't all have the same proposition set, and we certainly don't seem to place them all in the same order of importance. This is an issue about weights and clusters of points in the space, and is reserved for future discussion at this point.

New propositions can be introduced by refining old ones. For example, if we both understand "green", we can qualify "green" into subcategories of the taxonomy by

green.1 Green like grass.
green.2 Green like waterweed.

and if we want to subclassify green.1

Green.1.1 Green like Kentucky Bluegrass.
Green.1.2 Green like fescue grass.

And I understand there are agricultural tribes who have 100 different subcategories of green. It's only worth subclassifying green, instead of saying "Green like grass" if you want to say "Green like grass" a lot.

What green like grass means is that subset of the cluster of instances of green whose Hamming distance is closest to the cluster of grasses. But to talk about Hamming distances here one needs perhaps the methods of spectroscopy and the triangular coordinate system to describe hues.

Otherwise one just says "green like grass" or "green like waterweed" and leaves it at that.

People in the same specialty profession for example have a whole universe of discourse not accessible to those outside. There's a whole lot of issues about using weights to discriminate between categories which I have not thought thru properly yet. But an example is a recent discussion between Bruce Nevins and Bill Cunningham where Bruce had two meanings of agreement - one essentially an agreement about something in the past, and the other a commitment to make something happen in the future. I am inclined when being nitpickish to call the second kind a contract.

Level one introduces a new instance in the universe of discourse, essentially by giving its' Hamming distance from other instances, e.g.

"there's this kind of hallucination like an elephant but pink"

i.e. it has the elephantine Aristotelian attributes except for color pink, and being a perception that may not have been experienced directly by the recipient of the message. But the recipient can either say "I know what you mean" conveying at least a possibility of no directly shared experience, but probably having heard others talk about pink elephants, or he/she can say "Yes I saw one of those early on Jan 1st. 1992" conveying having drunk not wisely but too well the night before.

Level two is the confirmation by the recipient that a match has been found in recipient's Kanerva database.

I think higher levels exist and have to do with weights or values, and clusters that are perceived as similar by some community.

For example, when I fight with my boss, I am using a value system where scientific merit is highest in the list of weights, so I can perceive small differences in scientific merit. He poor soul is wrestling with the fact that if expenditure exceeds income by more than epsilon, then our whole operation will close down. So he, quite reasonably is more concerned about whether there is a market than about whether the idea is new and earthshaking. No wonder we disagree. The fact that I perceive his dilemma does not make me any happier about using the last few years when I can still think reasonably clearly to clean dull Augean stables.

All this is not well thought through. I think there are some interesting distinctions though. Some Aristotelian attributes are selectors for a universe of discourse, others are discriminators within the universe.

If one has a set of axioms relating Aristotelian attributes, neither set of attributes is clearly unique, unless one has no redundancy or non orthogonality. I think the formal statement of this is that if one falsifies an attribute, this should not lead to inconsistency in the model (in the sense of first order logic). But all kinds of snakes spring from this - did Boole REALLY find the "Laws of Thought", and what about the non standard logics. Avery - Aren't there some people in the Dept. of Philosophy at Canberra who worry about non standard logics?

Well, I'm rambling. I could perhaps have made this neater, but it would not have touched on as many of the questions I find interesting had I done so.

John Gabriel (gabriel@eid.anl.gov)

Date: Tue Nov 17, 1992 10:27 pm PST
Subject: *CSGintro* revision

[From Dag Forssell (921117.22.00)]

For Gary Cziko, Bill Silvert and others.

Here is a reformatted version of CSGINTRO.DOC with a max of 72 characters per line, as suggested by Bill Silvert. Len Lansky suggested 64, then reversed himself, so 72 it is. I shall make a point of posting @ 72, now that I know.

Gary: When I spoke with Mary, she mentioned that it would be nice to include your paper in the literature list. I second Mary's motion. Can you add it please. Your paper is dynamite.

But be sure you don't mess up MY document again. |:-)

Bill, this is a minor update. Please substitute if not much trouble.

INTRODUCTION TO THE CONTROL SYSTEMS GROUP NETWORK (CSGnet)
AND TO THE CONTROL SYSTEMS GROUP
Prepared by Dag Forssell 921117.

This introduction provides information about:

Why you might want to join the CSG and read CSGnet
Our subject: The control paradigm and Perceptual Control Theory
The evolution of the control paradigm
Demonstrating the Phenomenon of Perceptual Control
The purpose of CSGnet
CSGnet participants
How to ask effective questions
The Control Systems Group
Subscribing to CSGnet
How to obtain text and program files
Literature references

WHY YOU MIGHT READ THE CSGnet

If you have an interest in how people and organisms work.
If you are dissatisfied with the explanations in many of the life
sciences and would like a more rigorous approach....
If you insist on thinking things through for yourself....

OUR SUBJECT: THE CONTROL PARADIGM AND PERCEPTUAL CONTROL THEORY

Human control is the primary subject of CSGnet, but all forms of control
are game. Here are two introductions by Bill and Mary Powers:

There have been two paradigms in the behavioral sciences since the
1600's. One was the idea that events impinging on organisms make them
behave as they do. The other, which was developed in the 1930's, is
PERCEPTUAL CONTROL THEORY (PCT). Perceptual Control Theory explains
how organisms control what happens to them. This means all organisms
from the amoeba to humankind. It explains why one organism can't
control another without physical violence. It explains why people
deprived of any major part of their ability to control soon become
dysfunctional, lose interest in life, pine away and die. It explains
what a goal is, how goals relate to action, how action affects
perceptions and how perceptions define the reality in which we live
and move and have our being. Perceptual Control Theory is the first
scientific theory that can handle all these phenomena within a
single, testable concept of how living systems work.

William T. Powers, November 3, 1991

PERCEPTUAL CONTROL THEORY

While the existence of control mechanisms and processes (such as
feedback) in living systems is generally recognized, the implications
of control organization go far beyond what is generally accepted. We
believe that a fundamental characteristic of organisms is their

ability to control; that they are, in fact, living control systems. To distinguish this approach from others using some version of control theory but forcing it to fit conventional approaches, we call ours Perceptual Control Theory, or PCT.

PCT requires a major shift in thinking from the traditional approach: that what is controlled is not behavior, but perception. Modelling behavior as a dependent variable, as a response to stimuli, provides no explanation for the phenomenon of achieving consistent ends through varying means, and requires an extensive use of statistics to achieve modest (to the point of meaningless) correlations. Attempts to model behavior as planned and computed output can be demonstrated to require levels of precise calculation that are unobtainable in a physical system, and impossible in a real environment that is changing from one moment to the next. The PCT model views behavior as the means by which a perceived state of affairs is brought to and maintained at a reference state. This approach provides a physically plausible explanation for the consistency of outcomes and the variability of means.

The PCT model has been used to simulate phenomena as diverse as bacterial chemotaxis, tracking a target, and behavior in crowds. In its elaborated form, a hierarchy of perceptual control systems (HPCT), it has lent itself to a computer simulation of tracking, including learning to track, and to new approaches to education, management, and psychotherapy.

Control systems are not new in the life sciences. However, numerous misapprehensions exist, passed down from what was learned about control theory by non-engineers 40 or 50 years ago without further reference to newer developments or correction of initial misunderstandings. References in the literature to the desirability of positive feedback and the assertion that systems with feedback are slower than S-R systems are simply false, and concerns about stability are unfounded.

The primary barrier to the adoption of PCT concepts is the belief - or hope - that control theory can simply be absorbed into the mainstream life sciences without disturbing the status quo. It is very hard to believe that one's training and life work, and that of one's mentors, and their mentors, must be fundamentally revised. Therefore, PCT appeals to those who feel some dissatisfaction with the status quo, or who are attracted to the idea of a generative model with broad application throughout the life sciences (plus AI and robotics). There are very few people working in PCT research. Much of its promise is still simply promise, and it meets resistance from all sides. It is frustrating but also tremendously exciting to be a part of the group who believe that they are participating in the birth of a true science of life.

Mary Powers, November 1992

The PCT paradigm originates in 1927, when an engineer named Harold Black invented the negative feedback amplifier, which is a control device. This invention led to the development of purposeful machines. Purposeful machines have built-in intent to achieve consistent ends by variable means under changing conditions.

The discovery and formalization of the phenomenon of control is the first alternative to the cause-effect perspective ever proposed in any science.

The first discussion of purposeful machines and people came in 1943 in a paper called: Behavior, Purpose and Teleology by Rosenblueth, Wiener and Bigelow. This paper also argued that purpose belongs in science as a real phenomenon in the present. Purpose does not mean that somehow the future influences the present.

The first specific suggestion on how to use the concept of control to understand people came in 1957 in a paper entitled: A General Feedback Theory of Human Behavior by McFarland, Powers and Clark.

In 1973 William T. (Bill) Powers published a seminal book called "Behavior: the Control of Perception," which still is the major reference for PCT. See literature below.

This book spells out a complete model of how the human brain and nervous system works like a living perceptual control system. Our brain can be viewed as a system that controls its own perceptions. This view suggests explanations for many previously mysterious aspects of how people interact with their world.

Perceptual Control Theory has been accepted by independently thinking psychologists, scientists and other interested people. The result is that an association has been formed (the Control System Group), several books published, this CSGnet set up and that at latest count 16 professors are teaching PCT in American universities today.

DEMONSTRATING THE PHENOMENON OF CONTROL

The phenomenon of control is largely unrecognized in science today. It is not well understood in important aspects even by many control engineers. Yet the phenomenon of control, when it is recognized and understood, provides a powerful enhancement to scientific perspectives.

It is essential to recognize that this phenomenon exists and deserves an explanation before any of the discourse on CSGnet will make sense.

Please download the introductory demonstration demla.exe, which is an interactive program and/or rubberbd.txt, which is a text telling you how to demonstrate the phenomenon to yourself and a friend using only two rubber bands.

THE PURPOSE OF CSGnet:

CSGnet provides a forum for development, use and testing of PCT.

CSGnet PARTICIPANTS

Many interests and backgrounds are represented here. Psychology, Sociology, Linguistics, Artificial Intelligence, Robotics, Social Work, Neurology, Modeling and Testing. All are represented and discussed. A quality of participants on this net is that most are prepared to question and reconsider what they think they know.

HOW TO ASK EFFECTIVE QUESTIONS

Since PCT puts much conventional, well established wisdom on its head, it is helpful to begin by demonstrating the phenomenon of control to yourself and studying a few references. As you catch on, read this net and follow a thread that interests you for a month or more.

When you ask a question, please consider that in order to give you a good answer, a respondent will need to put your question in context.

Therefore, please introduce yourself with a statement of your professional interests and background. It will help if you spell out which demonstrations, introductory papers and references you have taken the time to digest and what you learned.

People on this net are in various stages of learning and understanding PCT. When you get a reply to your post, please consider that the respondent who found your question of interest and invested time in a reply, may benefit from knowing how you perceived the answer. Did it answer your question? Was it clear? Were you able to understand it?

THE CONTROL SYSTEMS GROUP

The CSG is an organization of people in the behavioral, social, and life sciences who see the potential in PCT for increased understanding in their own fields and for the unification of diverse and fragmented specialties.

Annual dues are \$45 for full members and \$5 for students (subsidized).

An annual meeting is held in Durango, Colorado, on the campus of Fort Lewis College. In 1993 it will begin in the evening of Wednesday, July 28, and end Sunday morning, Aug, 1. There will be 7 plenary meetings (mornings and evenings), with afternoons, mealtimes, and late night free for further discussion or recreation. Full details will be available on the net or by mail after April 1, 1993.

Net subscribers find it useful to have thematic collections of some of the network discussions, and it enables non-net members to keep up with them. Threads from this net are published on a quarterly basis in a booklet called the Closed Loop. These booklets are distributed to members and are available separately. A complimentary copy of Closed

Loop will be sent upon request. Back issues are available: Volume 1 (4 issues) is \$12. Single issues of Volume 2, beginning with Jan. 1992, are \$6 each.

For membership information and back issues of Closed Loop, write: CSG, c/o Mary Powers, 73 Ridge Place CR 510, Durango, CO 81301-8136.

SUBSCRIBING TO CSGnet

When you subscribe to CSGnet, you get this message, CSGINTRO.DOC. But you may have received it from a friend who printed it, seen it on a demodisk, or seen it on Usenet. To subscribe, send a message as follows: (Internet address followed by two message commands, one per line)

```
LISTSERV@VMD.CSO.UIUC.EDU
Subscribe CSG-L Lastname, Firstname, Affiliation, City, State.
help
```

(Lastname, Firstname, etc is optional commentary, but helpful).
("help" requests a list of most commonly used commands).
(The Bitnet address is: LISERSERV@UIUCVMD).
(This server is not sensitive to upper or lower case letters).

CSGnet can also be accessed via Usenet (NetNews) where it is listed as the newsgroup bit.listserv.csg-l.

HOW TO OBTAIN TEXT AND PROGRAM FILES

A number of ASCII documents and binary computer programs are available on a fileserver maintained by Bill Silvert. It is possible to download all these files via e-mail. If you are on internet, it is easiest to obtain binary program files via anonymous FTP. If you are on MCI mail, you have read about how you can transfer binary files with Kermit or Zmodem protocols. (Type help at the MCI mail prompt for directions). But the server cannot send binary files over the internet mail network, so download uee.scr first, then request the binary files uuencoded as ASCII files. The Internet address for the server is BIOME.BIO.NS.CA. CSGnet files are kept in the subdirectory pub/csg.

To get basic information and a current listing of available documents, send a message as follows: (Internet)

```
To:          SERVER@BIOME.BIO.NS.CA.
```

```
Commands:   help
            ftp
            get csg/Index
            end
```

"help" requests commands and explanations.
"ftp" asks details on anonymous FTP for internet.
"get csg/Index" requests the Index for the csg subdirectory.

Pay attention to letter case for commands! DOS is not dos.

As part of the index (of the csg directory), you may be looking at:

programs/msdos:

dem1a.exe 128437 Bill Power's demonstration of perceptual control
dem2a.exe 123649 Bill Power's modelling of control

documents/forssell:

uud.scr 53406 ASCII Compile uud.exe w DOS debug. Directions @ end.

If you want dem1a.exe (uuencoded) to get a "live" demonstration of the phenomenon of control, and the ASCII file uud.scr with directions at the end on how to use DOS debug to compile uud.exe to decode it, send the following message commands:

```
uue csg/programs/msdos/dem1a.exe
get csg/programs/forssell/uud.scr
```

The uuencoded dem1a.exe will be sent in four parts. Remove headers and use an editor to make it into one file (starting with table and ending with end) before you use uud.exe to restore the file. dem1a.exe is a self-extracting archive file. Put it in it's own directory before you execute it. You get complete documentation and a running program.

LITERATURE REFERENCES

For a complete list of CSG-related publications, get the file biblio.pct from the fileserver as described above. Here are some selected books and papers on Perceptual Control Theory:

Powers, William T. (1973). Behavior: The control of perception.
Hawthorne, NY: Aldine DeGruyter.
The basic text.

*Robertson, Richard J. and Powers, William T. (Eds.). (1990).
Introduction to modern psychology: The control theory view.
Gravel Switch, KY: CSG Book. (\$25 postpaid)
College-level text.

*Powers, William T. (1989). Living control systems: Selected papers.
Gravel Switch, KY: CSG Books. (\$16.50 postpaid)
Previously published papers, 1960-1988.

*Powers, William T. (1992). Living control systems II: Selected papers.
Gravel Switch, KY: CSG Books. (\$22 postpaid)
Previously unpublished papers, 1959-1990.

*Marken, Richard S. (1992). Mind readings : Experimental studies of purpose. Gravel Switch, KY: CSG Books. (\$18 postpaid)
Research papers exploring control.

Marken, Richard S. (Ed.). (1990). Purposeful Behavior : The control theory approach. American Behavioral Scientist, 34(1). (Thousand

Oaks, CA: Sage Publications)

11 articles on control theory.

Runkel, Philip J. (1990). _Casting nets and testing specimens_. New York: Praeger.

When statistics are appropriate; when models are required.

Hershberger, Wayne. (Ed.). (1989). _Volitional action: Conation and control_ (Advances in Psychology No. 62). New York: North-Holland.

25 articles (not all PCT)

Ford, Edward E. (1989). _Freedom from stress_. Scottsdale AZ: Brandt Publishing.

A self-help book. PCT in a counseling framework.

Gibbons, Hugh. (1990). _The death of Jeffrey Stapleton: Exploring the way lawyers think_. Concord, NH: Franklin Pierce Law Center.

A text for law students using control theory.

McClelland, Kent. (1992). _Perceptual control and sociological theory_.

Not yet published. Available from the author,
Grinnell University, Grinnell, Iowa.

McPhail, Clark. (1990). _The myth of the madding crowd_. New York: Aldine de Gruyter.

Introduces control theory as a way to explain group behavior.

McPhail, Clark., Powers, William T., & Tucker, Charles W. (1992). Simulating individual and collective action In temporary gatherings. _Social Science Computer Review_, _10_(1), 1-28.

Computer simulation of control systems in groups.

Petrie, Hugh G. (1981). _The dilemma of inquiry and learning_. Chicago: University of Chicago Press.

Introduces PCT to educational theory.

Richardson, George P. (1991). _Feedback thought in social science and systems theory_. Philadelphia: University of Pennsylvania Press.

A review of systems thinking, including PCT.

*These items are available from CSG Books, 460 Black Lick Road,
Gravel Switch, KY, 40328.

Date: Wed Nov 18, 1992 4:39 am PST

Subject: analog NN chip

Communications and Digital Signal Processing (CDSP)

Center for Research and Graduate Studies

CDSP Seminar

Title: The Integrated Neurocomputing Architecture (INCA)

Speaker: Mark Dzwonczyk, The C. S. Draper Laboratory

Here are a few notes for those of you who wonder what you missed. There were two issues of interest here, one regarding analog hardware for neural nets and one regarding sonar signal processing.

Apparently analog hardware implementations of neural nets are in disrepute because Intel tried and failed to make it work. JPL has tried again, however, this time successfully. This talk was about a proof-of-concept project to apply JPL's chip to several practical problems. The value of an analog implementation is that it takes 11 transistors per neuron, whereas a digital implementation takes more like 6000 transistors. Thus an analog implementation is better when miniaturization and low power requirements are important.

What Draper did was put several of JPL's chips onto a VME board in a fixed configuration: a 4-layer feedforward neural net with 64 nodes per layer, or 256 nodes altogether. Training was done offline on a Sun Workstation, and the final connection weights were "ftp'd" to the hardware. To train they simply applied backprop using commercial-off-the-shelf software (NeuralWare). (They are now working on hardware-in-the-loop training.) When asked how long training took, he said simply, "three months of a graduate student's time."

There were several applications that they used to "prove the concept," and the one he talked most of was realtime sidescan sonar target detection. The data looks like distance vs. time map, with pixels colored according to the amplitude of the echo. Targets are white (loud) blips lost in a bunch of noise, and it can be hard for a person to see them. The algorithm was to slide a 10x10 window around the image and look for underwater mines. Using 2 hidden layers, they achieved a false alarm rate of .03%. It is impressive that this straightforward approach to the problem yielded such a high degree of accuracy. (Nevertheless, the navy wants to cut this false alarm rate by another factor of 10 before fielding it, because the penalty for being wrong about a mine is rather high.)

jeff morrill

Subject: Handbook of Intelligent Control

From: Donald Sofge <sofge@ai.mit.edu>
Date: Mon, 16 Nov 92 19:25:25 EST
To: connectionists@cs.cmu.edu
Subject: Book Announcement

<<<<----- New Book Announcement ----->>>>

HANDBOOK OF INTELLIGENT CONTROL:
Neural, Fuzzy, and Adaptive Approaches

Edited by David A. White and Donald A. Sofge

Handbook of Intelligent Control provides a hands-on approach to integrating intelligent control approaches with existing control architectures for the nonlinear control of complex multivariate systems. It is an attempt by leading industry and academic researchers to present the current state-of-the-art developments in intelligent control in a highly readable, often tutorial format such that many of the techniques described within may readily implemented by the reader. The goals of this approach are:

- % To provide firm mathematical and theoretical foundations for current intelligent control methods
- % To demonstrate how these methods may be effectively combined

with existing control practices

% To provide overviews and extensive bibliographic references to much current work in this field

% To provide examples of real-world applications in robotics, aerospace, chemical engineering, and manufacturing which have served as a driving force behind new innovations in intelligent control.

Discussions concerning these applications are provided for two main reasons: to demonstrate how existing intelligent control techniques may be applied to real-world problems, and to provide challenging real-world problems which serve as impetus for new innovations in intelligent control and for which intelligent control solutions may provide the only solutions.

This book is an outgrowth of three major workshops held under the National Science Foundation Intelligent Control Initiative. The chapters included in this volume are implementations and extensions of creative new ideas discussed at these workshops. The application and integration of neural, fuzzy, and adaptive methods into "real-world" engineering problems makes this an ideal book for practicing engineers, as well as graduate and academic researchers. This volume contains the following chapters:

Foreword - Paul Werbos, Elbert Marsh, Kishan Baheti, Maria Burka, and Howard Moraff

Editors' Preface Donald Sofge and David White

Introduction to Intelligent Control

- 1 Intelligent Control: An Overview and Evaluation
Karl Strom and Thomas McAvoy
- 2 An Introduction to Connectionist Learning Control Systems
Walter Baker and Jay Farrell
- 3 Neurocontrol and Supervised Learning: An Overview and Evaluation
Paul Werbos

Conventional Control and Intelligent Approaches

- 4 Fuzzy Logic in Control Engineering
Reza Langari, Hamid Berenji, and Lotfi Zadeh
- 5 Adaptive Control of Dynamical Systems Using Neural Networks
K.S. Narendra
- 6 Optimal Control: A Foundation for Intelligent Control
David White and Michael Jordan
- 7 Development and Application of CMAC Neural Network-based Control
Gordan Kraft, Tom Miller, and D. Dietz

Applications of Intelligent Control

- 8 Artificial Neural Networks in Manufacturing and Process Control
Judy Franklin and David White
- 9 Applied Learning-Optimal Control for Manufacturing

Donald Sofge and David White

- 10 Neural Networks, System Identification, and Control in the Chemical Process Industries
Paul Werbos, Thomas McAvoy, and Ted Su
- 11 Flight, Propulsion, and Thermal Control of Advanced Aircraft and Hypersonic Vehicles
David White, Albion Bowers, Ken Iliff, Greg Noffz, Mark Gonda, and John Menousek

Advances in System Identification, Optimization and Learning Control Theory

- 12 Reinforcement Learning and Adaptive Critic Methods
Andrew Barto
- 13 Approximate Dynamic Programming for Real-time Control and Neural Modeling
Paul Werbos
- 14 The Role of Exploration in Learning Control
Sebastian Thrun

Publication Date: September 1992
568 pages Price: \$59.95

Available from: Van Nostrand Reinhold
VNR Order Processing
P.O. Box 668
Florence, Kentucky 41022-0668
Call Toll-free: 1 (800) 926-2665

-----END OF FORWARDED MESSAGES-----

This put me in mind of what the review section of the Journal of Living Control Systems might look like.

Bruce bn@bbn.com

Date: Wed Nov 18, 1992 6:01 am PST
Subject: formatting line length

[From: Bruce Nevin (Wed 921118 08:39:13)]

Another consideration: a 65-char line is 6.5 inches on typical ASCII display or printout, pretty standard for text. (Two inches for side margins on 8.5 x 11 paper.) Longer lines get harder to read. Since 65 (or maybe 64) is the limit for Cincinnati, perhaps we should all set our window sizes, editor margins, whatever, to format text accordingly. If you just have to insert returns by hand, a mark at the top and bottom of your screen can support a reference perception. Other than that, those who will be bothered by overlong lines will have to work out some kind of accommodation with those who can't be bothered

to tailor them. With drop-ins and new entrants this will be a recurrent issue.

The above text fits the constraint.

Of such processes are born bureaucratic regulations and other things perceived as social control, as witness reaction of one (no longer with us) to the convention about date stamps at the beginnings of csg-1 posts.

Bruce bn@bbn.com

Date: Wed Nov 18, 1992 6:14 am PST
Subject: Complexes of perceptions

[From Bill Powers (921117.2000)]

Note to Len Lansky: just set the margins in your word processor wider -- say, to 79.

Martin Taylor, Rick Marken --

The reason I use words like intensity...system concept is to avoid getting embroiled in loose language, including terms like "standard," "value", "belief," and so on.

Martin, you say you would like to use "belief" to mean a structured perception. I think we agree that what you mean by a structured perception is a collection of scalar perceptual signals. But a collection of signals, just by coexisting, does not create any special kind of perception. I feel my right big toe against the floor and I feel my left forefinger hit the "f" key. That does not create a structured perception.

What creates something structured out of a collection of scalar perceptual signals is a higher level perceptual function. If I make the toe-sensation at the same time as the finger-sensation I perceive simultaneity, a perception of relationship. The effect of this higher level perceptual function is to represent the presence of a particular kind of structuring AS A SIGNAL, in this case a signal standing for the presence of a relationship called simultaneity. But the presence of simultaneity means nothing unless it is perceived and explicitly represented.

The basic tenet of HPCT is that nothing is perceived unless it is represented as a perceptual signal. This includes what you call "structuring" of a set of signals. If a structuring is perceived, that is only because there is a perceptual function organized to represent a collection of inputs in terms of that kind of structuring.

This is just one aspect of a much deeper point I have been trying to make since BCP, where I brought it up as Rosenblatt's Principle. I have spoken about it on the net as the difference between implicit and explicit functions. In any collection of n perceptions, there is an infinity of implicit functions of those perceptions. That is, you could construct an arbitrary function of, for example, the form $ax_1 + bx_2 + \dots$ or any other function, assign arbitrary values to the coefficients, and continuously compute that function of the perceptions x_n . Actually doing that by using a computing device that outputs a signal continuously representing the value of that function makes the function explicit. To give that function of the variables any effect in a model, the model must contain the machinery for computing it. It is not enough for a bystander to know that such a function COULD be computed; that has no effect in the model. The function must

actually be computed if it is to have any meaning in the model, or any effects on anything else in the model.

In some neural-network models I have seen, perceptual processes are treated as the existence of organized patterns throughout some volume of the brain. The problem is that those patterns are being recognized by the wrong brain, the brain of the observer. All active networks carry signals that are exhibiting some pattern; which pattern depends on what regularities the observer is prepared to perceive. There might be some patterns that the observer finds especially easy to see -- synchronism and so forth. The catch is that these patterns simply represent the internal working of some part of the brain, but do not necessarily have any significance to any other part of the brain. For a pattern to have significance elsewhere in the brain, something must RECOGNIZE the pattern and indicate its presence or state by a signal sent to other parts of the brain. In short, there must be a perceptual function that converts one of the infinity of possible implicit patterns into an explicit signal. There is no such thing as "the" pattern of signals.

For a brain modeler to speak of any pattern of signals in a network, that modeler must be using a human perceptual function capable of indicating presence of that pattern. To the human observer, there can be just a hint of the pattern, a pretty clear sense of the pattern, or a perfect prototype pattern. There is a scale of perception ranging from hardly any to the maximum possible. So in the human observer, the pattern is represented as a scalar perception.

So why does the human modeler not put that same sort of function into the model? I claim it is because most human modelers see no connection between the system they are modeling and themselves. They reify their own perceptual interpretations. If they see a certain pattern in the behavior of a model, they assume that the pattern actually exists in the model, and that by its very existence it can have specific effects. They do not realize that their own pattern-recognition (or better, - construction) abilities are creating a significance in their own brains in the form of a perceptual signal. Yet they believe that somehow that collection of signals in the model, without any function to recognize it, can have significance in the model just by being in the model.

This is a very difficult point to communicate. I hope it's getting across, at least a little. A "complex of perceptions" appears as something unitary only to a perceptual function capable of perceiving that kind of thing and signalling that it is present. Two different perceptual functions receiving the same set of lower-order perceptions can easily perceive different "complexes" in it. The complex is in the perceptual function of the beholder.

Best, Bill P.

Date: Wed Nov 18, 1992 9:07 am PST
Subject: code metaphor

[From: Bruce Nevin (Wed 921118 10:35:06)]

The problem with using words like proposition and belief in new PCT senses that have no connection with natural language or formal "language" is that they are misleading to

readers who are consciously or unconsciously committed to the conduit metaphor of communication that has prevailed since the stoics and especially since Locke. In this metaphor, a message or belief inside the sender is translated or encoded into a code of sounds or marks (writing) and then translated or decoded back into a message or belief inside the receiver. The obvious alignment of your Layered Protocol theory of communication with this metaphor, which enhances its palatability to most audiences, gives me severe intellectual indigestion. Our discussions over the past almost two years have not made it clear to me where you stand on the literal merit of the metaphor.

To see that this metaphor is wrong you have to recognize that a code has no information structure in it other than that imparted to it by the process of encoding; that is, on this view an utterance can exhibit only such information structure as was previously present in the message or belief held to be "in" the sender.

There is no warrant for claiming that perceptions associated with an utterance (as the meanings of the utterance) are structured in the way that the utterance is, and there is much to suggest that they are not.

Suppose that we were to try to uphold such a claim. If the structure in utterances follows from prior structure in perceptions to be conveyed by them as "messages," then we could reasonably expect languages to be alike to the extent that people's respective universes of perception are alike. We would furthermore expect that agreement about words would constitute agreement about the associated perceptions. But the structural likenesses across languages are mostly formal matters not bearing strongly on differences of meaning or information, and the differences across languages with respect to semantic ranges of words and the like are striking and seem to far exceed what plausibly might be differences in the perceptual hierarchies (apart from language) of their speakers. Furthermore, while we might expect the information structures found in utterances to change over time or vary from one community to another (viz. for participants in a scientific subdiscipline with advancements in the science), we would not expect archaisms of language to persist despite changes in the organization of perception (as they do), and we would not expect change and variation in language to be independent of changes in perceptions or "messages" being "encoded" (as they frequently are).

Perceptions are structured by the organization of the perceptual hierarchy and we presume by structure that really does exist in the environment. Some of this is structure in a social environment: if people went away it would be no more, and if the participants in the particular community went away or even if they were replaced by people who had learned (sufficiently) differently, it would likewise be no more. The structure that is in utterances heard by a child learning language was previously learned, in a comparable environment, by those participating in it and thereby perpetuating it.

There is good reason to suppose that the global information structure of a discourse (including conversation, Penni) is a reflection of the organization of nonverbal perceptions, but the linearization of that structure in a particular discourse embodies constraints to which the essentially parallel nature of perceptual control is not subject. Even preferred ways of linearizing are socially institutionalized. We prefer a temporal metaphor for story telling ("and then what happened?"). People in certain other cultures find this incomprehensible and prefer a metaphor of spatial arrangement, and such cultures differ in the background pattern presumed for this arrangement. We find such discourse meandering and incoherent. Other cultures have developed yet other conventions, relatively little is known as study of this has only recently been taken up in earnest, and indigenous cultures are being wiped out at a rate of a dozen or more a year.

It seems probable that words and some multi-word constructions (and their reductions) are correlated with variously ramifying networks of nonverbal perceptions in the human perceptual control system, perhaps at the category level as Bill has proposed. But the particular correlations, and changes in them, are learned and negotiated socially, they are not byproducts of perceptual control. It does not take much exploration of bilingual dictionaries even of languages so closely related as French and English to see that words taken as basic translations have different ranges. That the differences in category-level perception that presumably follow have insufficient behavioral consequences to support a strong form of the Whorf-Sapir hypothesis is I think due to the fact that behavioral consequences are constrained by perceptual control quite independently of the conventions of language. This too indicates that language cannot be a code, since the structure that is in utterances cannot be an encoding of prior structure in perceptions taken as a "message."

I suggest that it is a worthwhile challenge to formulate a brief statement of your Layered Protocol Theory of communication in PCT terms without recourse to any vestige of the message/code/conduit metaphor. I think it will be stronger for the exercise. I certainly will be much happier with it.

Bruce bn@bbn.com

Date: Wed Nov 18, 1992 10:02 am PST
Subject: words & images;logical perceptions; intro

[From Bill Powers (921118.0730)]

Avery Andrews (921118.1430) --

>>When I hear "large tree at the top of a hill next to the path" I
>>construct something like an imagined top of a hill with an
>>imagined big tree next to the path.

>This won't work, because images have the wrong kind of information
>in them. A classical example from philosophy: an image of a
>triangle has to either be isosceles or not isosceles, but the concept
>triangle is neither.

You must understand that when I say "image" I mean a perceptual signal derived from inputs and indicating that something is present in those input signals, as per PCT. If I tell you "draw a triangle" you just draw a triangle -- it doesn't matter whether it's equilateral or isosceles or obtuse, or whether the end of the last line exactly touches the start of the first without crossing it. Yet you always draw a SPECIFIC triangle. Any triangle will give rise to a perceptual signal indicating that triangleness is present. If you want a special kind of triangle, you have to specify TWO reference signals: triangleness and equilateralness, for example. Then you are asking for perceptions of both triangleness and equilateralness, which are independent dimensions of variation and can be controlled (more or less) independently. If I ask you to draw an equilateral figure, you might draw a triangle, a square, a hexagon, an octagon, and so forth. You might make some of the vertices in more complex figures indented. In the uncontrolled dimensions, anything will do. All that matters in controlling for equilateralness is that all the sides be perceived as equal.

The same is true for the Ronald Reagan image. To control for such an image, all that matters is that some aspect of RRness be present, sufficient for a perceptual function to

wake up and signal "I see my guy." Cartoonists know how to wake up such perceptual functions by picking out a sufficient set of inputs and exaggerating them to make up for distortions of other inputs. That's called a caricature.

>It is indeed something of a mystery how anything can translate between
>images and the `concepts' that we like to represent as propositions in
>logic, but that mystery has got to be solved anyway, & once it is, this
>solution could be used to build a concept-to-image comparator.

In HPCT there are no such things as concepts. We break down perceptions in a different way into lower level and higher level perceptions of specific hierarchically-related classes (maybe not the right classes, but it's a start). In the informal language of pre-HPCT philosophy there are only a few categories of perception and they are poorly defined: you have concrete and abstract, or sensations-perceptions-concepts, and other such vague notions.

Propositions are generally composed of symbols related by implicit operations. In "Socrates is mortal" the operation is that of equality or interchangeability or class membership, depending on how you think of it. "Socrates" is the name of a configuration, a person. "is" is the name of a relationship or a class membership. "Mortal" is the name of a rather fuzzy collection of perceptions having to do with dying (which for oneself can only be imagined), time duration, and perhaps some principles like "Life ends."

But in dealing with "Socrates is mortal" as a proposition, the actual meanings of the terms make no difference; the words are just place-holders. The proposition is "x is a member of class y." It is then put together with other similar propositions at the logic level, or the ninth level in current HPCT, to yield a perception of TRUE or FALSE. If a conclusion is TRUE, and the reference level is TRUE, then the logical control process is satisfied and the logical process stops. The output proposition is then in the state that causes the input symbols to be in states that make the truth value of the perceived proposition match the reference value.

Logic has its uses, and its abuses. It helps in trying to figure out complex relationships. But it can be used to do logical manipulations of symbols where the symbols stand for perceptions that don't actually follow the logical rules. You can start by saying "All men are bicycles; Socrates is a man; therefore Socrates is a bicycle" and nothing at the logic level will object. Logical and other rule-driven processes may be powerful and complex, but they are stupid. They are certainly not the highest level of function in the brain. If they were, we would see nothing wrong with concluding that Socrates is a bicycle.

My chief objection to overmathematicized approaches to modeling is that mathematics doesn't need to make sense. There has to be a higher perception of the appropriateness of the mathematics. The judgement as to whether or not to believe the mathematical conclusions is not itself a mathematical process. It involves principles and system concepts -- the aspects of perception and action that we call, in older fuzzier terms, "wisdom."

John Gabriel (921117.2131) --

See my last paragraph to Avery.

>My baseline concept is a Kanerva database of truth values for a set
>of N Aristotelian propositions that together discriminate within an
>area of discourse. The truth values are bit strings N long, and

>can be thought of as being in a space of N dimensions, containing
>2**N points each representing a possible object in the area of
>discourse. Typical values of N are in the range of 100-200, and the
>space is very sparsely populated.

As no reference to meaning has been made so far, this is a strictly logic-level proposition. The universe of discourse is being DEFINED to have the desired properties: the rules of the game are being laid out.

>New propositions can be introduced by refining old ones. For >example, if we both understand "green", we can qualify "green" into >subcategories of the taxonomy by

green.1 Green like grass.
green.2 Green like waterweed.

In real perceptions, there are also

green.3 Green like about 2/3 of the way from grass to
waterweed.
green.4 Green like about 4/7 of the way from grass to
waterweed.

....

By assigning meanings to symbols as you propose, you're introducing an arbitrary grid into the world of lower-level perceptions, such that perceptions can be symbolized accurately only at the intersections of the grid lines, and jump from one intersection to another instantaneously (as, for example, the grass dries out). In the world of symbols there are no states of green between green.1 and green.2. Therefore if one person is perceiving such an intermediate shade of green, it will be impossible (save by luck) to convey that shade of green to the other person because the only terms available for description refer to fixed points in the continuum.

This means that the world of discourse represents the experienced world by forcing it into discrete forms. This limits the meanings that can be conveyed, and that can be perceived to have been conveyed. The fewer the agreed-on Aristotelian propositions that are used, the more stereotyped the discourse becomes. Even in the limit, allowing the maximum conceivable number of agreed-on propositions, the world of direct experience is always far more detailed than what we can say about it.

Dag Forssell (921117.2200) --

Finally I'm getting around to your intro to CSGnet.

I'm of two minds about this document. I know that if I were a stranger to the net, I would be put off by it. I don't need to be told why I might want to read CSGnet, particularly not when the telling comes across as a sales pitch. You say I'll be interested "If you're interested in how people and organisms work... If you're dissatisfied with explanations ... if you insist on thinking things through for yourself...". Of course that means that if I don't want to read CSGnet I must be some sort of stupid dolt who isn't interested in people or organisms, who is satisfied with nonrigorous explanation, and who doesn't like to think for himself. This sounds like an insurance salesman who says "Do you care what happens to your children after you die?"

I wrote my intro thing at your house and was being a salesman. I'm uncomfortable now when I read it.

Mary's stuff looks fine, but I have to say that, don't I? Let others judge.

In general, I draw back from the air of the expert aiding the neophyte who is wandering through CSGnet in a state of mild confusion. Most of the people who look in on CSGnet will be smarter than we are, and have long experience with approaching new ideas and learning about them efficiently. I don't think they need to be told how to ask effective questions. I think a description of the available materials should be enough -- nobody is going to consult them just because you tell them to, and most people who browse around networks know enough to read the source material before starting to express opinions.

I think you're doing us a service by working up this document, but I also hope it can be revised some more to eliminate the salesmanship and condescension.

Bruce Nevin (921118) --

The analogue neural net circuit sounds interesting. The transistor count per neuron is about the same as for an operational amplifier. Any idea where I should write to get the specs?

Best to all, Bill P.

Date: Wed Nov 18, 1992 11:43 am PST
Subject: Re: dear Rick

M. Lee,
Your dissertation sounds interesting but complex. I may be quite interested in what you have to say about knowledge transfer from group to individual, though, because my dissertation has to do with what kinds of information people select from their group's texts in order to in turn write to their group in their own, new texts. This board helps me keep focused on the idea that an individual can only select what s/he perceives, and perceives only what s/he is oriented to perceive...I think.

Cindy Cochran Depts. of English U. of Illinois / Carnegie Mellon U.

Date: Wed Nov 18, 1992 11:43 am PST
Subject: "beliefs", standards

[From Rick Marken (921118.0900)] Martin Taylor (921117 16:40]

>I have used "belief" to represent a structured perception.

Bill Powers (921117.2000) said EXACTLY what I planned to say about this (only he said it better) but he got to it first. Excellent post, Bill.

Martin Taylor (921117 17:00) --

>I'm a bit confused. I knew this was foolhardy territory to get into, but
>I can't see what the discussion of interaction procedures that evolve into
>conventions known as standards has to do with the system level.

This is part of the definition problem. The talk about standards is highly ambiguous -- sometimes I think people are talking about reference levels and sometimes about perceptual variables. I thought we had clarified it earlier a bit -- my conclusion was that Ed Ford (the main "standards" guy) uses the word standards to refer to "higher level perceptual variables" -- types of perceptions that might be described by words like

system concept, principle, value, belief, etc. He tries to help people set the "right" reference levels for these perceptions. So my reply to your post was really aimed at Ed -- I just don't think control of higher order variables is any more important in social interactions than control of other perceptual variables.

>Standards include greeting patterns, dress codes, thank-you notes for gifts,
>and all sorts of things for which the external appearance is what matters.

So what you mean by standards is "perceptual variables that involve another person". Well, now we have another possible meaning for "standard". Why don't we just stick to the PCT model terminology (and semantics)?

>As you can see, I don't think standards are anything that people control

Boy, you got me. In the quote above it sounded like standards were social perceptions. Now they are something that can't be controlled. And yet people reorganize when controlling them produces conflict. So it must not be failure to control standards that is leading to reorganization. But the reorganization leads to new, stable standards. So standards are a perceptible (to Martin) side effect of reorganizing to control perceptions that are not standards? In other words, people control perceptual variables; this can appear to an observer as a process of converging on social standards. Is this it? If so, I completely agree.

Best regards Rick

Date: Wed Nov 18, 1992 11:57 am PST
Subject: patter

[From: Bruce Nevin (Wed 921118 12:07:22)] (penni sibun 921113.1300) --

You described two situations, which you subsequently characterized as party patter and as flirting. You said that in party patter, little information was communicated though language was much used, and in flirting much information was communicated though language was little used.

My response was that in the party patter in fact considerable linguistic information is in the language used, but that it is not information that the participants care much about; and in the flirtation little information is in fact transmitted, however much is imagined, and that the most important thing communicated is interest in and availability for negotiating mutual agreements that have a more explicit foundation.

>in my description of flirting, i gave specific instances of
>information being conveyed (eg, ``i like it when you do that'').

You are imagining that you have conveyed the information that is in the sentence "I like it when you do that." He may very well be imagining something else, such as "she thought that was funny". Over time you can rule out a general emotional tone in what he is imagining (e.g. if he were imagining "she's making fun of me again" the flirtation would soon come to an end), and you can assume the general tone of "I'm enjoying your company", but probably nothing even as specific as your really rather vague "I like it when you do that" (e.g. doing it again would be boring on the second or third iteration and repulsive soon after). Even the more specific "I liked it when you just did that" has this problem.

There is information here, to be sure, but it is information about relationship of the same sort that all mammals communicate nonverbally. Nonverbal, gestural communication is more reliable than language in its proper domain, relationships and attitudes, but the kinds of information that are constituted in utterances are quite beyond its capacities. Conversely, this is why language is good for error-free transmission of linguistic information, but not particularly good for communication, in the sense of communication that happens in flirtation and is restrained in party patter.

Bruce bn@bbn.com

Date: Wed Nov 18, 1992 12:00 pm PST
Subject: the image

>(Avery Andrews (921118.1430)) --

> images have the wrong kind of information in them.
>A classical example from philosophy: an image of a triangle has to
>either be isocoles or not isocoles, but the concept triangle is neither.
>It's got to be something that subsumes both. Or another example, an
>image of Ronald Reagan on a horse, unobscured, will have the horse
>white or not white, Reagan with or without a hat, etc., while if I
>say `walk down the corridor until you get to the picture of RR on a
>horse', neither of these attributes matters.
>
>It is indeed something of a mystery how anything can translate between
>images and the `concepts' that we like to represent as propositions in
>logic, but that mystery has got to be solved anyway, & once it is, this
>solution could be used to build a concept-to-image comparator.

Bill's proposal covers this I think, about the input function of a category-level ICS accepting a disjunction of input signals corresponding to various perceptions, including a word or words. Thus, your imagined perception of RR on a white horse with a hat satisfies the category perception but without precluding other possibilities or even bogging much when encountering another instantiation. One could even go down the hall without a specific image in mind, but searching for any satisfier (other than the words "picture of RR on a horse"--aye, there's a rub).

Bruce bn@bbn.com

Date: Wed Nov 18, 1992 1:19 pm PST
Subject: *CSGintro* revision -Reply

[len lansky (921118.13110)]

The text did not all come through because of "exceeded mhs length limit" here. It stopped at paragraph "This book spells out a complete model of how the human brain and system....."

Would you be kind enough, Dag, to send me the remainder. I want to use it for a graduate seminar in social psychology on control theory. I will be asking for more help in planning same in a few weeks--the seminar is for winter quarter. Thanks

len lansky

Date: Wed Nov 18, 1992 1:20 pm PST
Subject: Re: code metaphor

[Martin Taylor 921118 13:30] (Bruce Nevin 921118 10:35:06)

>There is no warrant for claiming that perceptions associated with
>an utterance (as the meanings of the utterance) are structured in
>the way that the utterance is, and there is much to suggest that
>they are not.

In the Layered Protocol Theory (LP) the matter of structure is entirely internal to the protocol concerned. In LP there is no warrant for relating the structure of any virtual message to that of another type. They are essentially incommensurate.

Having said that, I will back off a bit, and point out a construct we call "global coherence". Inasmuch as lower-level virtual messages support the passage of higher-level ones, there will seem to be a sequential structure if an analyst looks only at the lower-level ones. But this structure actually is the province of the higher-level protocol, and is, so far as the lower one is concerned, only in the mind of the analyst.

>To see that this metaphor is wrong you have to recognize that a
>code has no information structure in it other than that imparted
>to it by the process of encoding; that is, on this view an
>utterance can exhibit only such information structure as was
>previously present in the message or belief held to be "in" the
>sender.

Technically, if I understand what you mean, that is incorrect. But I may be misunderstanding you, so bear with me. In a strictly coding procedure (no metaphor), the code may be redundant, and usually is. Therefore the code itself contains information relating to whether a particular received pattern was likely to have been the one that was sent. This kind of coding redundancy is only possible if the encoder and the decoder contain matched models of the coding procedure, so that the decoder can determine the probabilities for all possible messages sent by the encoder.

Convention in communication provides to only possibility for encoding to occur, and I identify the existence of encoding with syntactic structure. In LP syntax is not expected to be important in face-to-face conversation for levels much above the phrase, but it is important in going from acoustics to phonetics, and phonetics to syllable and probably word. It is what makes speech recognition machines possible when they have no contextual understanding of what the speech is for.

> If the
>structure in utterances follows from prior structure in
>perceptions to be conveyed by them as "messages," then we could
>reasonably expect languages to be alike to the extent that
>people's respective universes of perception are alike.

One would not expect this at levels below, perhaps the elements of an argument, and probably not even there if there are conventions with respect to how arguments are conducted. The argument is backwards. If there is encoding, it is arbitrary and must have nothing to do with what is encoded, and therefore one would expect the languages of

the world to be as different as the temporal segregation of their communities would allow. But I would, on the other hand, argue that the physical constraints of communication, based on informational efficiency, constrain the kinds of differences one could find among languages. I would think it unlikely that the parallelism inherent in perception would be found in a kind of language that introduced a small element of every item to be communicated at the beginning of an utterance, and added further small aspects of every item in each small time increment, until they all finished together. I would expect all languages to show some pattern of contiguity in time among items that are conceived to belong together in perception. After all, it is important that a warning about a tiger should be imparted quickly, whereas the location of a nice bunch of berries can wait.

>the particular correlations, and changes in them, are learned and
>negotiated socially, they are not byproducts of perceptual
>control.

I'd argue that the first part of this sentence is correct, the second not. See yesterday's discussion with Rick about "standards." I think linguistic and other social conventions are developed in exactly the same way, by reorganization that leads to minimum perceptual error overall.

>I suggest that it is a worthwhile challenge to formulate a brief
>statement of your Layered Protocol Theory of communication in PCT
>terms without recourse to any vestige of the message/code/conduit
>metaphor. I think it will be stronger for the exercise. I
>certainly will be much happier with it.

I won't do that, because, as I said, I do believe that there are social conventions that do allow (demand) specific encoding practices, especially at the lower protocol levels. But if you (or anyone) is interested, I now have annotated the slides I used for my Paris talk, and they, together with the tutorial, provide an introduction to the way I see LP theory being linked to PCT. They may not satisfy you, and they are certainly not complete, but they are an introduction to both PCT and LP as well as the link between them (50 pages tutorial, 70 pages annotated slides, so don't ask unless you are really interested--I will send to Bruce without him asking).

Martin

Date: Wed Nov 18, 1992 1:21 pm PST
Subject: Comments on Promoting PCT

[From Ray Jackson (921118)] for Dag Forssell

Dag, a few comments on your Promoting PCT post (921116.1940)

>UNDERSTANDING A THEORY-BASED LEADERSHIP PROGRAM

I was thinking maybe using the title to ask the question: WHY A

THEORY-BASED LEADERSHIP PROGRAM? That's what most would be asking.

>Summary

The summary is GREAT. All I would suggest is that you emphasize that the education your program offers provides a focus for the consistent ends achieved through the varying means of human behavior.

>1. PCT is an acronym for Perceptual Control Theory.

Instead of a footnote from the TOC, would you want to name the theory and explain the acronym in the summary?

>BOLD: The value of a good theory

>The power and practicality of a good theory is awesome. It also

>provides clarity.

Here would be a good spot to comment on how theories offer us a framework for what we understand...that it Rbolsters our common senseS When I'm working with people, I suggest that thinking without a theory (or paradigm), is like merely talking about cities... and that the theory (or paradigm) provides us with a map to place them.

>...As we shall see, the history of science recounts many phenomena

>that were considered mysterious and too complicated until an

>appropriate concept was offered.

Excellent. Most executives and managers I've worked with seem to regard performance technology as magical, and success with the human element as a matter of luck (although they don't openly admit it). In fact, I'm convinced that my PCT-based interpretations of behavioral phenomena at my plant tend to be minimized because Ray is a teacher, or the kind of guy who Rknows peopleS; I have some kind of a "gift".

>A good theory makes real breakthroughs possible.

But I'm still not sure the reader knows why. Maybe the piece on Kuhn and paradigms needs to be worked in here.

>BOLD: Objective "REALITY" versus personal "reality"

I really like this description of reality...in fact, I used it the other night when Tom Hancock let me introduce his Ed Research class to PCT (by the way, he's adopted Runkel's book Q Casting Nets Q as the text).

>BOLD: Individual development

>...This is why the power of a good theory is awesome.

It's beginning to come clear to me.

>BOLD: Existing training programs

>Companies spend millions of dollars on training.

The hot button; maybe you could add to the emphasis of the statement by asking the reader to consider the true effectiveness of these spent dollars.

>To illustrate the variety, this list...

Nice touch, most managers are very familiar with this stuff in particular. Also, you may want to add that a major focal point of most new training programs is that our program is THE ONE...but they never say why it's different from the rest.

>In most cases, people have fun and like the training, but four or
>five months later, little has changed in the workplace. I believe a
>large reason for this is that the training is situational /
>anecdotal and focuses on "what/how to DO." Each participant is left
>to integrate the many disparate lessons of the training experience
>into the framework of their personal reality, such as it is.

I love it. It's too bad this is business as usual for most personnel and training programs.

>People ask: Show me what to DO (cause) so I will get results
>(effect). Our program shows clearly why this is fallacious, but
>that does not change the fact that this is what many people have
>come to expect.

This is a key interest of mine. Not only is this what people have come to expect, but this is what they want. Traditionally, education and training has been a two-dimensional cause-effect world; it's outstanding that PCT offers the third dimension of consistent results through varying means. Certainly, this focus enables me and other teachers to be better educators, but there is still the dilemma of the cause-effect mindset towards training in most learners. More on this in a later post.

>Instructions on "what/how to DO" are valid only in a given set of
>circumstances. Typically a training scenario is carefully selected
>and told with drama and humor by a speaker. You are told what the
>circumstances were, what was done and what the results were. You
>imagine that the same thing will happen if you do the same thing.
>You feel euphoric as you imagine success. A large part of the
>"what/how to DO" training does not really apply in individual cases
>because the world is full of varying conditions and changing
>disturbances. The lessons become irrelevant and are soon forgotten.
>Euphoria fades away.

The guys who own CareerTrack probably hate you (that's alright, they're awfully rich by now). Don't resent them too much, they will probably ask you to do a video soon.

>BOLD: Theory-based education and training

>PCT explains a wide variety of phenomena of everyday experience.

One thing you may want to emphasize is WHY theory-based education is more effective, which you could then reiterate the practical aspects of in the following section.

>BOLD: Speed, cost, effectiveness

>Instead of using *multiple* programs, each one covering some aspect

>of human interaction, you can use *one* to understand yourself and
>others in some detail...

One thing you may want to mention is the adaptability of theory-based education (so the reader doesn't actually think you use the SAME curriculum in every situation).

>BOLD: Scientific revolutions

I love Kuhn's material, and I find a high level of interest every time I present it; what about using one of the great quotes from Kuhn himself?

>BOLD: One example of a scientific revolution

Good example, but you end up with two, don't you (including Newton)? By the way, has Galileo been in purgatory all these few centuries, or would he have actually been sent to hell; and when the church acknowledged the error, was he able to enter heaven? Oh never mind, I guess that's for another Network.

>BOLD: Perceptual Control Theory

There is considerable challenge in presenting a brief description of PCT to those with a traditional mindset. Since I've been struggling with a this, I found Mary Power's intro posted earlier on the net very helpful. So helpful, in fact, I used it (stole, copied, plagiarized...) extensively in the name of PCT. Seriously, I was able to modify it to a single-page summary which I've used with managers and students alike. (Mary, thank you very much!).

The summary:

Perceptual Control Theory: a brief summary.

While the existence of control mechanisms and processes (such as feedback) in living systems is acknowledged, the implications of individuals as control systems goes far beyond what is usually considered in terms of human behavior or motivation. Control Theorists believe that the fundamental characteristic of organisms is their ability to control; that they are, in fact, living control systems, constantly interacting with their environment to maintain a desired state. This view is referred to as Perceptual Control Theory, or PCT.

PCT requires a major shift in psychological thinking from the traditional approach: that is, what is controlled is not behavior, but the perception of a current reality. The traditional behavioristic view is a two-dimensional cause/effect model, which considers behavior as a dependent variable, producing a finite response to external stimuli. However, this approach provides no explanation for the phenomenon of achieving consistent ends through varying means. Attempts to model behavior as planned and computed output that can be demonstrated to the point of precise calculation are unobtainable in a real environment that is changing from one moment to the next.

The PCT model views behavior as the means by which a perceived state of affairs is brought to and maintained at a reference state. This approach provides a physically plausible explanation for a third dimension in human behavior: the consistency of outcomes by the variability of means. The PCT model has been successfully used to explain phenomena as diverse as bacterial chemotaxis, tracking a target, piloting automobiles or airplanes, incidents in organizational behavior, and issues in human relationships. In its elaborated form, a Hierarchy of Perceptual Control Theory (HPCT),

has lent itself to new approaches to education, management, counseling and psychotherapy. As a prevailing paradigm for the depiction and explanation of human behavior, PCT holds much promise for establishing meaningful and accurate definitions of phenomena previously thought to be happenstance.

>The power of a program based on a good theory is awesome. >There is nothing as practical as a good theory based program. At this point, the reader knows why. You may want to emphasize that this is especially true with education and training programs.

The rest of the descriptions of who will pay attn (nice explanation of error signals), who can understand, what's in the program, and relevance to work and personal life are perfunctory and adequate. Anyone who's followed the paper to this point and/or is considering your program will find the information helpful (not to mention the PCT threads in the narrative).

Dag, I hope you find these comments encouraging and useful. I'm convinced that industry will eventually see the value of PCT when they realize that it is the essence of true empowerment. Keep at it.

Best Regards, Ray

Date: Wed Nov 18, 1992 1:30 pm PST
Subject: Re: Complexes of perceptions

[Martin Taylor 921118 11:30] (Bill Powers 921117.2000)

>>Martin, you say you would like to use "belief" to mean a structured
>perception. I think we agree that what you mean by a structured
>perception is a collection of scalar perceptual signals. But a
>collection of signals, just by coexisting, does not create any special
>kind of perception.
>...
>What creates something structured out of a collection of scalar
>perceptual signals is a higher level perceptual function.
>...
> The effect of this higher
>level perceptual function is to represent the presence of a particular
>kind of structuring AS A SIGNAL, in this case a signal standing for the
>presence of a relationship called simultaneity. But the presence of
>simultaneity means nothing unless it is perceived and explicitly
>represented.
>...
>The basic tenet of HPCT is that nothing is perceived unless it is
>represented as a perceptual signal. This includes what you call
>"structuring" of a set of signals.
>...
>This is a very difficult point to communicate. I hope it's getting
>across, at least a little. A "complex of perceptions" appears as
>something unitary only to a perceptual function capable of perceiving
>that kind of thing and signalling that it is present.

Well, I don't think it is a difficult point to get across, at all. It is the same point that AI researchers use in defining symbols to "mean" something, and the point against

which connectionists rail when they argue for distributed representations. In my discussion about structured percepts and structured control systems, I am taking an agnostic point of view on the matter. There is absolutely nothing about a structured percept that prevents that structure from being the result of a perceptual input function in a higher ECS, but neither is there any reason why such a higher ECS must exist.

HPCT does not have as a basic principle that the *conscious* perception of something depends on the existence of a scalar perceptual signal representing the CEV that the "something" corresponds to. HPCT says nothing explicitly about *conscious* perception (i.e. the everyday sense of the term "perception"). HPCT deals with the control of scalar perceptual signals, which in some as yet mysterious way presumably relate to conscious perception. Structured percepts are indeed in the mind of the analyst, but they are also composed of elements that are controlled perceptual signals, and the analyst sees that they tend to be controlled as a bunch--a complex. Quite conceivably, they are consciously perceived as a unitary multidimensional complex by the mind in which they reside, but that's not part of the deal.

No, your point isn't difficult. It is, I think, unnecessary. It raises the same flags as the "grandmother cell" does, and thereby warns people off directions of enquiry that might be profitable.

> Two different
>perceptual functions receiving the same set of lower-order perceptions
>can easily perceive different "complexes" in it. The complex is in the
>perceptual function of the beholder.

This is quite right, and there is no reason why a complex should not serve as input to any number of perceptual functions, while still acting as a complex. If you can imagine it, think of performing a golf swing (or a tennis racket). You can *consciously* perceive the swing, or the complex of muscle tensions, or both. This example doesn't exactly catch what I am talking about, because it refers to conscious perception, but it provides an analogy that might be useful.

Please note that I am remaining agnostic as to whether complexes as such have reality outside the mind of the analyst, or whether the higher ECSs that convert them to scalar perceptual signals are actually necessary.

Martin

Date: Wed Nov 18, 1992 3:04 pm PST
Subject: Re: "beliefs", standards

[Martin Taylor 921118 14:10] (Rick Marken 921118.0900)

>So what you mean by standards is "perceptual variables that involve another
>person". Well, now we have another possible meaning for "standard". Why
>don't we just stick to the PCT model terminology (and semantics)?

I've been trying to stick very precisely to the PCT model, but I don't know of any standard terminology to handle what we are talking about.

The problem with any definition of "standard" is that it is something (let's not say what) that one person applies to the observable actions of another. A person may apply

standards to themselves (isn't non-sexist language lovely?), but only as an observer, possibly in imagination, of their own actions.

Standards have a funny status. I cannot control your behaviour, because I have no sensory information that allows me to perceive it. But I can control my perception of your actions, in the same way I can control my perceptions of the inanimate world. To control perceptions of anything, I perform actions (not controlled; the product of all past reorganization), and if my reorganizations have been effective, my perceptual signals come closer to their reference levels. If I hold reference levels for my perceptions of your actions, the same applies. I act, and if the error signal does not decrease, there is a reasonable probability that I will reorganize. You, too. We both reorganize if controlling (as CEV's) each other's actions is not proceeding successfully (and if it matters--i.e if we are working at a reasonably high gain). Our mutual reorganization will probably wind up eventually in a situation where our perceptual errors are not too large. Then, each of us is acting according to the other's standards. This cannot happen if it causes a more than compensating increase in errors related to control of percepts outside the interaction. The most likely end result is that most people in a community use much the same set of standards.

Naturally, the end result of reorganizing through social interaction and the control of the actions (not the behaviours) of each other will be the existence of perceptual functions in each person that relate to patterns of actions in other people (and perhaps in themselves). Specific reference levels for these perceptions will be associated with the probability of low errors in other perceptual signals, and those reference levels may become the kind of "standards" that you were originally talking about.

>As you can see, I don't think standards are anything that people control

People don't control reorganization. Reorganization allows them to control their percepts more successfully.

>Boy, you got me. In the quote above it sounded like standards were
>social perceptions. Now they are something that can't be controlled. And
>yet people reorganize when controlling them produces conflict. So it must
>not be failure to control standards that is leading to reorganization.
>But the reorganization leads to new, stable standards.

Do you follow it better now? Think about what is happening in the interactions of many people, all potentially reorganizing in response to sustained error, rather than about the words used to describe the interactions. I don't mind changing terminology, but as you never tire of pointing out, PCT is a new and different viewpoint. To talk, we use old language that never was built on PCT, and take the old words into the PCT context with new, sharper meanings. I've done the same in LP, where I use "syntax" "semantics" and "pragmatics" in a precise sense that I think is included in the old vague senses. It can cause confusion, but short of creating a new language, what else can one do?

If, instead of "standard" we use "social convention," does this help?

> So standards are
>a perceptible (to Martin) side effect of reorganizing to control
>perceptions that are not standards? In other words, people control
>perceptual variables; this can appear to an observer as a process of
>converging on social standards. Is this it? If so, I completely agree.

I'm not sure this is what I said. We haven't arrived at a Paskian third-level agreement over an understanding, because I am not sure how to interpret this paragraph. Does the earlier part of this posting agree with it?

Martin

Date: Wed Nov 18, 1992 3:05 pm PST
Subject: Re: Internal/External

[Ray Allis 921118.1310

Over in one of the worlds next door, they were discussing "social construction" and associated topics . . .]

----- Begin Included Message -----

Paul Ceruzzi asks for elaboration on competing constructions of networks by telephone and computer people. This phenomenon is as old as networks themselves. Some of the earliest theoretical work on networking was done by Paul Baran of RAND (researched early 1960s, published 1964 under the title "On Distributed Communications.") When Baran had worked out his model of a digital communications network he went around to various people in the computer and communications industries to get their feedback. When he got to Bell Labs he was met with incredulity by the analog communications engineers, who were unable to believe that a reliable network could be created using digital switching (even though digital engineers at Bell Labs found Baran's ideas plausible). Engineers whose sense of what was possible was formed by experience in the analog world had a hard time putting their faith in ideas from the digital world, such as the idea that digital transmissions could be self-correcting. (This is not to imply that the analog engineers were somehow deficient. Their grasp of the needs and realities of telephone communications in general and AT&T in particular was probably acute). Baran's network was never built, apparently due to infighting among RANDS's Defense Department clients over who would get to build it.

The most obvious conflicts between telephone and computer people occur in the area of standardization. International networking standards are set by the CCITT, which is dominated by telecoms carriers. They tend to come up with standards that seem perfectly reasonable to telephone people but leave computer people perplexed and angry. I think the problem boils down to the two sides starting out with different assumptions about what is needed, leading to different constructions of what a "workable" standard is. For instance: telephone engineers expect steady, low-bandwidth traffic, while computer engineers expect the network to have to handle bursts of high-volume traffic. Telephone people think in terms of a continuous connection between two parties, whereas computer people see the transfer of data in discrete chunks as a desirable alternative. Etc. The result was that different systems were built by PTTs and computer people, and standardization is far from being achieved despite the existence of official "standards". (This topic is a large chunk of my dissertation; I can send you more details if you're interested.)

To get back to your original point: in order to understand what's going on in this situation, I needed to look at both the disputed technical details of the networks (to see what effect these would have on the functioning of the network) and the larger social/cultural situation (to understand why factors such as being a monopoly carrier or having concerns about data privacy might affect one's evaluation of the technology).

Janet Abbate

----- End Included Message -----

Date: Wed Nov 18, 1992 4:35 pm PST
Subject: "beliefs"

[From Rick Marken (921118.1300)]

Martin Taylor (921118 11:30) to Bill Powers (921117.2000)

> There is absolutely nothing about a structured percept that
>prevents that structure from being the result of a perceptual input function
>in a higher ECS, but neither is there any reason why such a higher ECS must
>exist.

Nothing except the way the model works.

>HPCT does not have as a basic principle that the *conscious* perception of
>something depends on the existence of a scalar perceptual signal representing
>the CEV that the "something" corresponds to.

Conscious/schmonscious -- a perception is a scalar perceptual signal in PCT.

>No, your point isn't difficult. It is, I think, unnecessary.

Well, I'm really glad that YOU know that YOU understand Bill's point. And thanks for helping me and Tom and Bill and the rest of us understand how the PCT model works.

> It raises
>the same flags as the "grandmother cell" does, and thereby warns people
>off directions of enquiry that might be profitable.

What is the problem with the "grandmother" cell? PCT says there are many grandmother cells -- for grandma intensities, sensations, transitions, configurations (this is probably the one you are thinking of), relationships, sequences, programs, principles, categories, system concepts. What's the "red flag"; what directions of inquiry are unprofitable -- the search for controlled perceptual variables, perhaps?

>This is quite right, and there is no reason why a complex should not serve
>as input to any number of perceptual functions, while still acting as a
>complex. If you can imagine it, think of performing a golf swing (or a tennis
>racket).

Forget imagining it; try modelling it.

>Please note that I am remaining agnostic as to whether complexes as such have
>reality outside the mind of the analyst, or whether the higher ECSs that
>convert them to scalar perceptual signals are actually necessary.

This is science, Martin; not theology. Show me the model that you have in mind; how do these "complexes" work in the PCT model?

Best regards Rick

Date: Wed Nov 18, 1992 6:21 pm PST
Subject: language

[Avery Andrews 921119.1136] (Bill Powers (921118.0730))

This helps a lot. I think people have to be very careful about terminology here, since I doubt that I'm the only person who is likely to react this way. `image' is a pretty controversial word these days. I think it would be best to reserve it for things like bitmaps, & find some other terminology for the higher-level perceptual signals.

(Martin Taylor (assorted postings))
presumably my perception of a social inhibition against checking for comprehension comes from spending most of my time in circumstances (e.g. Academia) where comprehension actually doesn't matter very much, and other values such as social cohesion prevail (e.g. it isn't helpful to make people think that you think that they're so dumb that they can't be expected to understand what you're saying).

Perhaps people whose perspective was formed in, say, naval aviation would have a different view of things (choosing that because, from what I've read about it, it seems to be the most complicated and dangerous thing that any group of people do on a routine basis).

Avery.Andrews@anu.edu.au

Date: Thu Nov 19, 1992 12:13 am PST
Subject: *CSGIntro* revised OK?

[From Dag Forssell (921118.23.00)]

For Gary Cziko, Len Lansky, Bill Powers and others.

Here is a reformatted, re-edited version of CSGINTRO.DOC with a max of 64 characters per line, as requested by Len Lansky and seconded by Bruce Nevin. I see no reason not to accomodate Cincinnati and others. The kilobytes stay the same.

Bill:

>I don't need to be told why I might want to read CSGnet,
>particularly not when the telling comes across as a sales
>pitch. You say I'll be interested "If you're interested in how
>people and organisms work... If you're dissatisfied with
>explanations ... if you insist on thinking things through for
>yourself...". Of course that means that if I don't want to read
>CSGnet I must be some sort of stupid dolt who isn't interested
>in people or organisms, who is satisfied with nonrigorous
>explanation, and who doesn't like to think for himself.

I see your point. It takes a long time to be sensitized to these things from old thinking. Now that you call my attention to it, it appears unnecessary anyway. It is gone. You have also influenced my next introduction letter. Thanks.

>I don't think they need to be told how to ask effective
>questions.

Fair enough, I have come to think that people would not learn courtesy from a statement like this anyway. I have cut 3/4.

>I wrote my intro thing at your house and was being a salesman.
>I'm uncomfortable now when I read it.

You are nuts. Here is a succinct statement of what a person can hope to learn from PCT. It is not condescending. Why should anyone be interested in PCT if you insist on telling them nothing, for fear of not being able to prove every implication 100.00%. We cannot shy away from telling people what's in it for them. What is wrong with being a salesman anyway? Are they not honorable people? Watch the condescension. I am proud to sell PCT. Is that dishonorable? Am I about to cheat people?

Let us not forget that "the competition" can prove NOTHING. PCT in its infancy is lighthouse years ahead of the massive volume of nonsense many smart people have studied for centuries. But condescension is out, and rightly so. People just don't know any better. I have not worked up to your incredible patience yet.

>I also hope it can be revised some more to eliminate the
>salesmanship and condescension.

Done?

>Mary's stuff looks fine, but I have to say that, don't I? Let
>others judge.

I would certainly like to hear from the others. I thought it was great. If this gets posted often enough, perhaps a few more people will notice and comment.

INTRODUCTION TO THE CONTROL SYSTEMS GROUP NETWORK (CSGnet)
AND TO THE CONTROL SYSTEMS GROUP
Prepared by Dag Forssell 921118.

This introduction provides information about:

- Our subject: Perceptual Control Theory
- The evolution of the control paradigm
- Demonstrating the Phenomenon of Control
- The purpose of CSGnet
- CSGnet participants
- Asking questions
- The Control Systems Group
- Subscribing to CSGnet
- How to obtain text and program files
- Literature references

OUR SUBJECT: PERCEPTUAL CONTROL THEORY

Here are two introductions by Bill and Mary Powers:

There have been two paradigms in the behavioral sciences since the 1600's. One was the idea that events impinging on organisms make them behave as they do. The other, which was developed in the 1930's, is PERCEPTUAL CONTROL THEORY (PCT). Perceptual Control Theory explains how organisms control what happens to them. This means all organisms from the amoeba to humankind. It explains why one organism can't control another without physical violence. It explains why people deprived of any major part of their ability to control soon become dysfunctional, lose interest in life, pine away and die. It explains what a goal is, how goals relate to action, how action affects perceptions and how perceptions define the reality in which we live and move and have our being. Perceptual Control Theory is the first scientific theory that can handle all these phenomena within a single, testable concept of how living systems work.

William T. Powers, November 3, 1991

PERCEPTUAL CONTROL THEORY

While the existence of control mechanisms and processes (such as feedback) in living systems is generally recognized, the implications of control organization go far beyond what is generally accepted. We believe that a fundamental characteristic of organisms is their ability to control; that they are, in fact, living control systems. To distinguish this approach from others using some version of control theory but forcing it to fit conventional approaches, we call ours Perceptual Control Theory, or PCT.

PCT requires a major shift in thinking from the traditional approach: that what is controlled is not behavior, but perception. Modelling behavior as a dependent variable, as a response to stimuli, provides no explanation for the phenomenon of achieving consistent ends through varying means, and requires an extensive use of statistics to achieve modest (to the point of meaningless) correlations. Attempts to model behavior as planned and computed output can be demonstrated to require levels of precise calculation that are unobtainable in a physical system, and impossible in a real environment that is changing from one moment to the next. The PCT model views behavior as the means by which a perceived state of affairs is brought to and maintained at a reference state. This approach provides a physically plausible explanation for the consistency of outcomes and the variability of means.

The PCT model has been used to simulate phenomena as diverse as bacterial chemotaxis, tracking a target, and behavior in crowds. In its elaborated form, a hierarchy of perceptual control systems (HPCT), it has lent itself to a computer simulation of tracking, including learning to track, and to

new approaches to education, management, and psychotherapy.

Control systems are not new in the life sciences. However, numerous misapprehensions exist, passed down from what was learned about control theory by non-engineers 40 or 50 years ago without further reference to newer developments or correction of initial misunderstandings. References in the literature to the desirability of positive feedback and the assertion that systems with feedback are slower than S-R systems are simply false, and concerns about stability are unfounded.

The primary barrier to the adoption of PCT concepts is the belief - or hope - that control theory can simply be absorbed into the mainstream life sciences without disturbing the status quo. It is very hard to believe that one's training and life work, and that of one's mentors, and their mentors, must be fundamentally revised. Therefore, PCT appeals to those who feel some dissatisfaction with the status quo, or who are attracted to the idea of a generative model with broad application throughout the life sciences (plus AI and robotics). There are very few people working in PCT research. Much of its promise is still simply promise, and it meets resistance from all sides. It is frustrating but also tremendously exciting to be a part of the group who believe that they are participating in the birth of a true science of life.

Mary Powers, November 1992

THE EVOLUTION OF THE CONTROL PARADIGM

The PCT paradigm originates in 1927, when an engineer named Harold Black invented the negative feedback amplifier, which is a control device. This invention led to the development of purposeful machines. Purposeful machines have built-in intent to achieve consistent ends by variable means under changing conditions.

The discovery and formalization of the phenomenon of control is the first alternative to the cause-effect perspective ever proposed in any science.

The first discussion of purposeful machines and people came in 1943 in a paper called: Behavior, Purpose and Teleology by Rosenblueth, Wiener and Bigelow. This paper also argued that purpose belongs in science as a real phenomenon in the present. Purpose does not mean that somehow the future influences the present.

The first specific suggestion on how to use the concept of control to understand people came in 1957 in a paper entitled: A General Feedback Theory of Human Behavior by McFarland, Powers and Clark.

In 1973 William T. (Bill) Powers published a seminal book called "Behavior: the Control of Perception," which still is the major reference for PCT. See literature below.

This book spells out a complete model of how the human brain and nervous system works like a living perceptual control system. Our brain can be viewed as a system that controls its own perceptions. This view suggests explanations for many previously mysterious aspects of how people interact with their world.

Perceptual Control Theory has been accepted by independently thinking psychologists, scientists and other interested people. The result is that an association has been formed (the Control System Group), several books published, this CSGnet set up and that at latest count 16 professors are teaching PCT in American universities today.

DEMONSTRATING THE PHENOMENON OF CONTROL

The phenomenon of control is largely unrecognized in science today. It is not well understood in important aspects even by many control engineers. Yet the phenomenon of control, when it is recognized and understood, provides a powerful enhancement to scientific perspectives.

It is essential to recognize that this phenomenon exists and deserves an explanation before any of the discourse on CSGnet will make sense.

Please download the introductory demonstration demla.exe, which is an interactive program and/or rubberbd.txt, which is a text telling you how to demonstrate the phenomenon to yourself and a friend using only two rubber bands.

THE PURPOSE OF CSGnet:

CSGnet provides a forum for development, use and testing of PCT.

CSGnet PARTICIPANTS

Many interests and backgrounds are represented here. Psychology, Sociology, Linguistics, Artificial Intelligence, Robotics, Social Work, Neurology, Modeling and Testing. All are represented and discussed.

ASKING QUESTIONS

Please introduce yourself with a statement of your professional interests and background. It will help someone answer if you spell out which demonstrations, introductory papers and references you have taken the time to digest.

THE CONTROL SYSTEMS GROUP

The CSG is an organization of people in the behavioral, social, and life sciences who see the potential in PCT for increased understanding in their own fields and for the unification of diverse and fragmented specialties.

Annual dues are \$45 for full members and \$5 for students (subsidized).

An annual meeting is held in Durango, Colorado, on the campus of Fort Lewis College. In 1993 it will begin in the evening of Wednesday, July 28, and end Sunday morning, Aug, 1. There will be 7 plenary meetings (mornings and evenings), with afternoons, mealtimes, and late night free for further discussion or recreation. Full details will be available on the net or by mail after April 1, 1993.

Net subscribers find it useful to have thematic collections of some of the network discussions, and it enables non-net members to keep up with them. Threads from this net are published on a quarterly basis in a booklet called the Closed Loop. These booklets are distributed to members and are available separately. A complimentary copy of Closed Loop will be sent upon request. Back issues are available: Volume 1 (4 issues) is \$12. Single issues of Volume 2, beginning with Jan. 1992, are \$6 each.

For membership information and back issues of Closed Loop, write: CSG, c/o Mary Powers, 73 Ridge Place CR 510, Durango, CO 81301-8136.

SUBSCRIBING TO CSGnet

When you subscribe to CSGnet, you get this message, CSGINTRO.DOC. But you may have received it from a friend who printed it, seen it on a demodisk, or seen it on Usenet. To subscribe, send a message as follows: (Internet address followed by two message commands, one per line)

LISTSERV@VMD.CSO.UIUC.EDU

Subscribe CSG-L Lastname, Firstname, Affiliation, City, State.
help

(Lastname, Firstname, etc is optional commentary, but helpful).
("help" requests a list of most commonly used commands).
(The Bitnet address is: `LISTSERV@UIUCVMD`).
(This server is not sensitive to upper or lower case letters).

CSGnet can also be accessed via Usenet (NetNews) where it is listed as the newsgroup `bit.listserv.csg-1`.

HOW TO OBTAIN TEXT AND PROGRAM FILES

A number of ASCII documents and binary computer programs are available on a fileserver maintained by Bill Silvert. It is possible to download all these files via e-mail. If you are on internet, it is easiest to obtain binary program files via anonymous FTP. If you are on MCI mail, you have read about how you can transfer binary files with Kermit or Zmodem protocols. (Type help at the MCI mail prompt for directions). But the server cannot send binary files over the internet mail network, so download uue.scr first, then request the binary files uuencoded as ASCII files. The Internet address for the server is BIOME.BIO.NS.CA. CSGnet files are kept in the subdirectory pub/csg.

To get basic information and a current listing of available documents, send a message as follows: (Internet)

To: SERVER@BIOME.BIO.NS.CA.

Commands: help
ftp
get csg/Index
end

"help" requests commands and explanations.
"ftp" asks details on anonymous FTP for internet.
"get csg/Index" requests the Index for the csg subdirectory.

Pay attention to letter case for commands! DOS is not dos.

As part of the index (of the csg directory), you may be looking at:

programs/msdos:

dem1a.exe 128437 Bill Power's demonstr of perceptual control
dem2a.exe 123649 Bill Power's modelling of control

documents/forssell:

uud.scr 53406 ASCII Compile uud.exe w DOS debug Dir @ end.

If you want dem1a.exe (uuencoded) to get a "live" demonstration of the phenomenon of control, and the ASCII file uud.scr with directions at the end on how to use DOS debug to compile uud.exe to decode it, send the following message commands:

uue csg/programs/msdos/dem1a.exe
get csg/programs/forssell/uud.scr

The uuencoded dem1a.exe will be sent in four parts. Remove headers and use an editor to make it into one file (starting with table and ending with end) before you use uud.exe to restore the file. dem1a.exe is a self-extracting archive file. Put it in it's own directory before you execute it. You get

complete documentation and a running program.

LITERATURE REFERENCES

For a complete list of CSG-related publications, get the file biblio.pct from the fileserver as described above. Here are some selected books and papers on Perceptual Control Theory:

Powers, William T. (1973). *Behavior: The control of perception*. Hawthorne, NY: Aldine DeGruyter.
The basic text.

*Robertson, Richard J. and Powers, William T. (Eds.). (1990). *Introduction to modern psychology: The control theory view*. Gravel Switch, KY: CSG Book. (\$25 postpaid)
College-level text.

*Powers, William T. (1989). *Living control systems: Selected papers*. Gravel Switch, KY: CSG Books. (\$16.50 postpaid)
Previously published papers, 1960-1988.

*Powers, William T. (1992). *Living control systems II: Selected papers*. Gravel Switch, KY: CSG Books. (\$22 postpaid)
Previously unpublished papers, 1959-1990.

*Marken, Richard S. (1992). *Mind readings: Experimental studies of purpose*. Gravel Switch, KY: CSG Books. (\$18 postpaid)
Research papers exploring control.

Marken, Richard S. (Ed.). (1990). *Purposeful Behavior: The control theory approach*. *American Behavioral Scientist*, 34(1). (Thousand Oaks, CA: Sage Publications)
11 articles on control theory.

Runkel, Philip J. (1990). *Casting nets and testing specimens*. New York: Praeger.
When statistics are appropriate; when models are required.

Hershberger, Wayne. (Ed.). (1989). *Volitional action: Conation and control* (Advances in Psychology No. 62). NY: North-Holland.
25 articles (not all PCT)

Ford, Edward E. (1989). *Freedom from stress*. Scottsdale AZ: Brandt Publishing.
A self-help book. PCT in a counseling framework.

Gibbons, Hugh. (1990). *The death of Jeffrey Stapleton: Exploring the way lawyers think*. Concord, NH: Franklin Pierce Law Center.
A text for law students using control theory.

McClelland, Kent. (1992). *Perceptual control and sociological theory*. Not yet published. Available from the author, Grinnell University, Grinnell, Iowa.

McPhail, Clark. (1990). The myth of the madding crowd.
New York: Aldine de Gruyter.

Introduces control theory to explain group behavior.

McPhail, Clark., Powers, William T., & Tucker, Charles W.
(1992). Simulating individual and collective action In temporary
gatherings. Social Science Computer Review, 10(1), 1-28.

Computer simulation of control systems in groups.

Petrie, Hugh G. (1981). The dilemma of inquiry and learning.
Chicago: University of Chicago Press.

Introduces PCT to educational theory.

Richardson, George P. (1991). Feedback thought in social
science and systems theory. Philadelphia: University of
Pennsylvania Press.

A review of systems thinking, including PCT.

*These items are available from CSG Books, 460 Black Lick Road,
Gravel Switch, KY, 40328.

Date: Thu Nov 19, 1992 3:39 am PST
Subject: Re: Complexes of perceptions

[From Oded Maler 921119]

[From Bill Powers (921117.2000)]

You seem to claim that a signal which is a combination of two
or more other signals "really" (explicitly) exists when there is a "gate" that takes the
two wires and outputs on a third wire a signal whose observed value is the combination of
the two.

Now consider a different situation where those two signals are not combined to create a
new perceptual function which is servoed, but rather are both servoed independently but
their effects are combined via the output function (e.g., two nerve fibers that feed into
neighboring parts of the same muscle).

In such situation a, there exist virtually a combined signal which can behave as-if it is
the controlled signal, but it does not really exist inside the system. Put it another
way, the "real" signals can be probed directly (in principle) by an external observer,
but in order to observe the "virtual" signals, the measurement instruments of the
observer should also include the missing circuits for creating the combinations.

--Oded

Date: Thu Nov 19, 1992 4:19 am PST
Subject: ADAPTIVE BEHAVIOR

From Greg Williams (921119)

Perhaps the following will be of interest to some netters. (Bill, if the arm paper is
rejected by SCIENCE, ADAPTIVE BEHAVIOR seems worth trying -- maybe more appropriate than
BIOLOGICAL CYBERNETICS.)

BEGIN INCLUDED MESSAGE

Alife Digest, Number 089
Monday, November 16th 1992

~~~~~  
~ Artificial Life Distribution List ~  
~  
~ All submissions for distribution to: alife@cognet.ucla.edu ~  
~ All list subscriber additions, deletions, or administrative details to: ~  
~ alife-request@cognet.ucla.edu ~  
~ All software, tech reports to Alife depository through ~  
~ anonymous ftp at ftp.cognet.ucla.edu in ~ftp/pub/alife (128.97.50.19) ~  
~  
~ List maintainers: Liane Gabora and Rob Collins ~  
~ Artificial Life Research Group, UCLA ~  
~  
~~~~~

[...]

Date: Tue, 3 Nov 92 16:17:27 +0100
From: meyer@biologie.ens.fr (Jean-Arcady MEYER)
Subject: Adaptive Behavior - Table of Contents

The first issue of Adaptive Behavior was released in August 1992. The second is under press.

For inquiries or paper submissions, please contact one of the editors:

- Editor-in-Chief: Jean-Arcady Meyer, France - meyer@wotan.ens.fr
- Associate Editors: Randall Beer, USA - beer@alpha.ces.cwru.edu
Lashon Booker, USA - booker@starbase.mitre.org
Jean-Louis Deneubourg, Belgium - sgoss@ulb.ac.be
Janet Halperin, Canada - janh@zoo.utoronto.ca
Pattie Maes, USA - pattie@media-lab.media.mit.edu
Herbert Roitblat, USA - roitblat@uhunix.uhcc.hawaii.edu
Ronald Williams, USA - rjw@corwin.ccs.northeastern.edu
Stewart Wilson, USA - wilson@smith.rowland.com

=====

ADAPTIVE BEHAVIOR 1:1
Table of Contents

A Model of Primate Visual-Motor Conditional Learning
by Andrew H. Fagg and Michael A. Arbib

Postponed Conditioning: Testing a Hypothesis about Synaptic Strengthening
by J. R. P. Halperin and D. W. Dunham

The Evolution of Strategies for Multi-agent Environments

By John J. Grefenstette

Evolving Dynamical Neural Networks for Adaptive Behavior
By Randall D. Beer and John C. Gallagher

=====

ADAPTIVE BEHAVIOR 1:2
Table of Contents

Adapted and Adaptive Properties in Neural Networks Responsible for Visual
Pattern Discrimination.

By J.-P. Ewert, T.W. Beneke, H. Buxbaum-Conradi, A. Dinges, S. Fingerling,
M. Glagow, E. Schurg-Pfeiffer and W.W. Schwippert.

Kinematic Model of a Stick Insect as an Example of a 6-legged Walking System.

By U. Muller-Wilm, J. Dean, H. Cruse, H.J. Weidemann, J. Eltze and F. Pfeiffer.

Evolution of Food Foraging Strategies for the Caribbean Anolis Lizard using
Genetic Programming.

By J.R. Koza, J.P. Rice and J. Roughgarden

Behavior-based Robot Navigation for Extended Domains.

By R.C. Arkin

END INCLUDED MESSAGE

Best wishes, Greg

P.S. Other lists don't seem to worry about line lengths of ca. 79 characters;
why should we? Caveat receiver, I say.

Date: Thu Nov 19, 1992 6:07 am PST
Subject: Conference: Complex Systems '92

----BEGINNING OF FORWARDED MESSAGES----

Complex Systems '92

FROM BIOLOGY TO COMPUTATION

The Inaugural Australian National Conference on Complex Systems

December 14-15, 1992

Australian National University
CANBERRA ACT

FINAL CIRCULAR

THE CONFERENCE THEME

Complex systems are systems dominated by non-linearity or interactions between their components. The last few years have seen an extraordinary growth of interest in complex systems. The underlying theme for this conference - from biology to computation - reflects trends that have been apparent in both biology and computer science.

On the one hand biologists have begun to accept the limitations of the reductionist approach and have started investigating emergent properties. One of the most potent paradigms linked with this new approach is the notion of life as a form of computation. This paradigm finds its ultimate expression in the new field of "artificial life".

On the other hand computer scientists, trying to discover how to organize the workings of massively parallel computers, quickly realized that they need to look at existing multiple processors. i.e. living organisms. This search, long apparent in the field of artificial intelligence, has focussed chiefly on neural networks. However, the lessons of biology are now being applied more and more widely, such as in the introduction of genetic algorithms.

The theme of this conference, then, reflects the dual notions of life as natural computation and computation as artificial life. The papers cover fields from artificial life to parallel computers including neural networks and social systems.

Please pass on this notice to interested colleagues. For further information contact the conference secretary or the convenors.

KEYNOTE SPEAKERS

Geoffrey Fox "Advances in parallel computing"

Professor Fox (Syracuse) is one of the world's leading authorities on parallel computers. At Caltech during the 1980's he pioneered the hypercube multiprocessor. He currently heads the Northeast Parallel Architecture Centre at Syracuse University.

Walter Freeman "Parallel computation in biological systems"

Professor Freeman (Berkeley) pioneered the study of complex behaviour in neural systems, with extensive published work in this field going back over two decades. His extensive biological studies are now yielding novel algorithms for machine pattern recognition.

Paulien Hogeweg "As large as life and twice as natural:
Bioinformatics and the artificial life paradigm"

Professor Hogeweg (Utrecht) pioneered the study of informatic processes in biotic systems. She has modelled self-organization in a vast range of biological systems, including prebiotic molecular evolution, cell signalling in the immune system, and the emergence of social behaviour in insect colonies.

Date: Thu Nov 19, 1992 6:11 am PST
Subject: Machine Learning List

A while back several people asked about my ref to a machine learning list. I just got another digest forwarded to me. The list address is ml@ics.uci.edu and the distributor (in the From: field of the email header) is Michael Pazzani, pazzani@ics.uci.edu (this parenthetical phrase is just to keep the period away from the address).

Bruce bn@bbn.com

Date: Thu Nov 19, 1992 9:11 am PST
Subject: *CSGintro* revised OK? -Reply

[len lansky (921119.0816)]

Thanks Dag. I got all of it and it will print just fine.

Thanks also to Bill Powers re idea about margins. I am working on it; the problem is getting rid of hard returns--I think I or our system is screwing up going back and forth between ascii and wordperfect. Thanks again.

Len Lansky

Date: Thu Nov 19, 1992 9:12 am PST
Subject: visual-linguistic-motor model

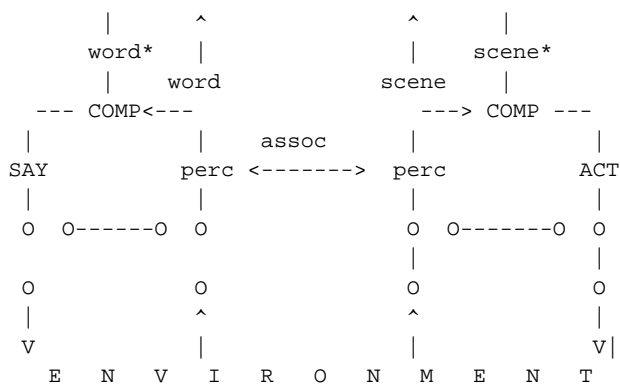
[From Bill Powers (921119.0700)]

RE: verbal & nonverbal control systems


I awoke this morning with a partly-formed idea about how words and nonverbal experiences get coordinated. The first thought was "Maybe Avery is right." It's possible to look for "a tree at the top of the hill next to the path" by continually looking at the scene and asking "Could that be called 'a tree at the top?'"

But it would also work if you had previously transformed the words into a reference perception in visual terms, and continually compared the current visual scene with the reference scene, in silence. My next thought was "Maybe some people do it one way and other people the other way."

Then I thought "Maybe we all can do it either way." That led to this diagram:



Note that the systems are drawn as mirror images to put the perceptual functions next to each other. The "O"s are switches: you fill in the gaps with a line, like this:



On the left:

If "SAY" is not connected to anything and
"perc" is not connected to anything

and on the right


"perc" is connected to the environment and
"ACT" is connected to the environment, or
both are connected in the imagination loop, then

acting to produce the reference scene in nonverbal perception on the right will produce the words associated with that scene in perception on the left. This is the connection shown above.

Clearly this would also work if we interchanged "word" and "scene". So this is a system that can be switched to allow the control of a nonverbal perception to evoke a corresponding verbal perception, or vice versa.

OK, put that half-finished idea aside. It's suggestive, anyway.

Then Avery's question about how you find the refrigerator floated to the surface. Right after it came this:



Note that I have used an acronym for a "perceptual control operator" which also happens to represent the Perceptual function, Comparator, and Output function in the correct relationship for our standard diagram. A perceptual control operator can be made of a sub-hierarchy of ECSs -- elementary control systems.

All the variables are now starred because they are reference variables. According to BPC, these are evoked from memory by address signals, shown as signals descending from Above. This addressing method of choosing reference signals allows a higher system operating in terms of one kind of perception to select a different kind of perception as a reference signal for a lower system, memory providing the translation from one logical type to another. An addressing dimension such as time would be common to all perceptual types. There could be others.

The horizontal double-headed arrows are suppose to indicate that all these stored perceptions are linked associatively. When any one of them is selected by a higher system, the others in this associative set are also selected. An associative set is formed by perceptions recorded with a common address -- say, time.

Now, suppose that you have looked in a direction X,Y and have seen a certain configuration centered in the visual field and have heard the word "refrigerator" spoken, all at the same time. This associative set, once recorded, now contains two directional perceptions, a visual configuration perception, and a verbal configuration perception.

If a higher system selects the X and Y angular coordinates as reference signals, it will also select the reference signal for a certain configuration, and another reference signal for a word. You will "expect" to see a particular familiar object and you will "expect" to hear it called "refrigerator." You will change your direction of looking until the sensed direction matches the selected reference direction X* and Y*, and you will see a refrigerator configuration matching the reference configuration, correcting that error. Not hearing the "refrigerator" configuration, you will imagine the word as you experienced it the first time, spoken; this will correct the error in the word-control system.

If a higher system selects the nonverbal configuration as a reference signal, it will also, by association, select an X and a Y coordinate, and will select a reference word. As above, you will look in the required direction, see the required configuration, and imagine the required word.

If a higher system selects the word-configuration "refrigerator" as the reference signal, you will perceive that word via imagination (if you don't say it aloud). At the same time, the visual refrigerator reference signal and a pair of direction reference signals X and Y will be selected associatively, so you will experience looking in the direction X,Y and experience the expected refrigerator configuration. The looking will correct the visual configuration error, as before.

Note that the configuration-PCO does not actually require any output function. Its error is corrected by the X and Y control systems which foveate the image of the refrigerator, thus supplying the requested perception of the refrigerator configuration.

Once you're looking at the refrigerator, it's a simple matter to increase its retinal size by facing in the direction of looking and moving forward, control a perception of a hand relative to the perceived handle as in the Little Man model, and use the arm control system to make the door perception (called "door") change to the open configuration (called "open"). Similar associative sets of reference signals would, I think, work here, too.

This arrangement requires that X and Y be defined in objective space, if you are to be able to locate the refrigerator from any position in the room and any initial orientation of the eyes, head, and body. Thus the PCO for direction control must contain a perceptual map of the room that is rotated and shifted according to the kinesthetically and visually sensed relative direction of looking. The relationship between the relative direction signals (eyes-head-shoulders-hips-knees-ankles-feet) and the room map must be updated often to keep the corrections in register. Wayne Hershberger, you may recognize elements of the oculomotor study here.

Likewise, the refrigerator perceptual function has to be able to provide a refrigerator perceptual signal with the refrigerator in a range of orientations relative to you. At the same time, the orientation perception of the refrigerator must remain in constant relationship to the orientation perception of the room, if the same scene is to be reproduced. These are all problems of configuration perception, which remain to be solved.

I think this conception could be turned into a model. It breaks down into several finite problems, which can be solved up to a point and from there on stipulated. We might, for example, decide that instead of angular directional reference signals we should have X*, Y*, and Z* Cartesian locations in the room, and leave the problem of converting from the sensed eyes-and-body configuration to objective space for later. We obviously don't need to figure out how words are recognized or produced; we can just say that they are. The associative part looks easy to implement in a limited situation.

I think we could make up a working model that could do the following: given the name of an object, turn and look at it; given a picture of an object, name it and turn to look at its prototype in the room; given a direction in which or a location at which to look, know that the object seen is a familiar object and name it. We would have to stipulate object recognition. Perhaps we could work in kinesthetic body configuration sensing, but we would have to stipulate, for the time being, conversion of visual images to objective space. But even with these limitations, I think this might be an interesting model. I invite anyone who is interested to be the first one to get such a model working.

Best to all, Bill P.

Date: Thu Nov 19, 1992 9:41 am PST
Subject: Re: Complexes of perceptions

[Martin Taylor 921119 11:00] (Oded Maler 921119)

>Now consider a different situation where those two signals are
>not combined to create a new perceptual function which is servoed,
>but rather are both servoed independently but their effects
>are combined via the output function (e.g., two nerve fibers that
>feed into neighboring parts of the same muscle).

Actually, one can go a little more microscopic than this. If it is true, as we discussed some months ago, that an ECS can be represented by a single neuron, then that neuron seems to do exactly the combination you are talking about. Each dendrite serves the reference-perceptual input-comparator functions, but they combine into a single output function. I really don't know what kind of control that permits for the perceptual signals derived from the separate dendrites. There is an intrinsic degree-of-freedom problem. Unless someone wants to try it out either mathematically or in simulation, it's probably best not to worry too much about this complication.

It isn't really what I was talking about in connection with distributed representation and the control of perceptual complexes, which I think is what you were responding to. Your point seems valid, but hard to analyze.

Martin

Date: Thu Nov 19, 1992 1:25 pm PST
Subject: Images, Dag's rewrite, virtual signals

[From Bill Powers (921119.1000)]

Misc public replies to private posts:

Martin, a purchase order for the Little Man v.2 will be fine. A Macintosh port of my demos by your people would be quite possible. Only library references would need changing. Will work with you as required.

Bruce, I would appreciate it if you could follow through on the neural net chip.

Rick Marken (921118.1300) --

It is important for a loose canon not to fire when pointing inboard.

Avery Andrews (921119.1136)--

On images:

The power of the pandemonium-type model really shows up when you talk about things like reference images. Neither a perception or a reference level having to do with an image needs to be a literal image (bitmap). The reference signal signals THAT the relevant dimension is present, and how much of it. To recall an image for perception via action or imagination one need only supply a nonzero scalar reference signal to the right control system. If lower-level systems are involved, you get the lower-order details of the image on the way to producing the "image- is-present" signal. If the imagination connection is short-circuited at the top level in question, you get the "image-is-present" signal without any corroborating details such as the image itself. You remember THAT the image was there, but you don't experience the image in terms of sensations, configurations, etc. How complete an image you get depends on how many dimensions of the perception you specify -- i.e., how many top-level control systems you send scalar reference signals to. How vivid the image depends on the level at which the imagination-connection is closed.

Dag Forssell (921118.23.00)--

Your rewrite captured exactly the tone I was hoping for. Very, very good! Suddenly the whole document, including the parts you didn't change, looks much better. As far as I'm concerned we can go with it.

I really admire your ability to act constructively on criticism. I even respect your opinion when you say I'm nuts (although I already knew that).

Oded Maler (921119)--

>Now consider a different situation where those two signals are not
>combined to create a new perceptual function which is servoed, but
>rather are both servoed independently but their effects are
>combined via the output function (e.g., two nerve fibers that feed
>into neighboring parts of the same muscle).

You're right about this. An excellent and even more extreme example is any of the so-called spinal "reflexes." The tendon reflex involves dozens or maybe hundreds of control systems acting in parallel, all their output signals contributing to the tension in a single muscle. When the muscle tenses, all the feedback signals rise at once. The net reference signal, averaged over all the comparators, leads to a set of error signals which again rise and fall together. The error signals enter the muscle, each one causing a different set of fibers to contract. All these effects finally come together into one single muscle tension that stretches the tendon. I spoke of this effect in BCP.

>In such a situation, there exist virtually a combined signal which

>can behave as-if it is the controlled signal, but it does not
>really exist inside the system.

Yes, we can represent all these parallel control systems as a single equivalent system containing single scalar signals. This is not a "complex" of signals like what Martin Taylor is talking about, because the effective signal is simply the sum of all the signals of the same kind (or something like it). Different patterns of signal intensity within each bundle of signals do not have any significance as far as the controlled variable or the action is concerned -- changing the distribution of signals within a bundle would not change the outcome if the sum were preserved.

>Put it another way, the "real" signals can be probed directly (in principle)
>by an external observer, but in order to observe the "virtual" signals, the
>measurement instruments of the observer should also include the missing
>circuits for creating the combinations.

Yes. In a way you're saying what I'm about to say to Martin in the next comment. You're pointing out that in order for the virtual signals to be known, the measuring instrument must contain a perceptual function that makes the virtual situation real by creating a meter reading. I'm saying that for the brain to make use of virtual signals it must do the same thing.

Martin Taylor (921118.1130) --

>Please note that I am remaining agnostic as to whether complexes as such
>have reality outside the mind of the analyst, or whether the higher ECSs
>that convert them to scalar perceptual signals are actually necessary.

I count agnosticism as progress.

My point is that a complex is not a unitary thing until it has passed through a perceptual function that determines what kind of unitary thing is there. If such a complex has reality in the mind of the analyst, that is evidence that the analyst has a perceptual function of the kind required to see it. If a human analyst can see unity in such a complex, then the model needs to have the ability to do the same thing. If the complex is left unperceived by this sort of function -- that is, if it is not represented by a scalar signal -- then it is just a collection. The difference is like that between a chemical mixture and a chemical compound -- except that something else has to do the compounding. None of this argues against complex perceptual functions in which signals interact before the final scalar outputs are produced. A set of perceptual functions could look like a perceptron, with all kinds of internal interactions, combinations, and feedbacks. But coming out of this perceptron will be individual output lines, each carrying a scalar signal that is a different function of the inputs to the perceptron. Each output line could then be represented as the output of a separate perceptual function, so the perceptron would be conceptually converted to a collection of input functions that all receive copies of the same set of inputs and independently compute their output perceptual signals. As I've said before, that's a convenient fiction, and we can't explain in that way how the perceptron actually works. The perceptron is the physically real organization. My representation just makes it easier to understand once it's organized.

By the way, has anyone ever thought of putting a constraint on a perceptron such that all its possible output signals are orthogonally related to the set of inputs? This would be extremely handy if we want one perceptron with multiple outputs to be part of a set of independent control systems, each controlling one of the perceptual outputs.

If I haven't made my mental image clear, what I'm thinking of is that a perceptron really represents a GROUP of input functions.

Best to all, Bill P.

Date: Thu Nov 19, 1992 2:28 pm PST
Subject: Re: Complexes of perceptions

[From Rick Marken (9921119.1000)]

In a rush to morning meeting but...

Oded Maler (921119) to Bill P.

>You seem to claim that a signal which is a combination of two
>or more other signals "really" (explicitly) exists when there
>is a "gate" that takes the two wires and outputs on a third wire
>a signal whose observed value is the combination of the two.

>Now consider a different situation where those two signals are
>not combined to create a new perceptual function which is servoed,
>but rather are both servoed independently but their effects
>are combined via the output function (e.g., two nerve fibers that
>feed into neighboring parts of the same muscle).

>In such situation a, there exist virtually a combined signal which
>can behave as-if it is the controlled signal, but it does not really
>exist inside the system. Put it another way, the "real" signals
>can be probed directly (in principle) by an external observer, but
>in order to observe the "virtual" signals, the measurement instruments
>of the observer should also include the missing circuits for creating
>the combinations.

From your description I get

$$o1 = f(p1) , o2 = f(p2)$$
$$o = g(o1 + o2)$$
$$p1 = h(o) ; p2 = h(o)$$

where p1,p2 are perceptual signals, o1 and o2 are their outputs and o is the muscle output resulting from the combined effect of o1 and o2.

I think you would find that p1 and p2 (the separate perceptual signals) are not independently controllable.

But maybe my representation of your suggestion is wrong? Could you help me out here, Oded.

It seems that you and Martin and maybe some others believe that perceptions can be represented in "distributed" groups of neurons. I don't understand the fascination with this latest trend-- I have no objection to it. I just don't understand how such a concept

is implemented in a model (such as pct) that USES these perceptual representations. In PCT it is true that perceptions (above level one) are implicitly represented in the vector state of a number of neurons - $p = f(x_1, x_2, \dots, x_n)$ where the x's are the states of the neuron's on which the perception depends. But the model, in order to use the perception implicit in the vector, computes a signal, p_i and the function, f , determines which of the many perceptions implicit in the vector is used in the model. My challenge to Martin (and now Oded) -- who appear to object to the scalar representation of perceptions in the PCT model -- is simply -- show me what YOU have in mind as a way to make "distributed" representations of perception usable by the PCT model.

I have nothing personal against "distributed" models of perception; just show me how such models can be used as part of a WORKING PCT model. Otherwise, such proposals sound more like ideologies than models.

Thanks Rick

Date: Thu Nov 19, 1992 3:26 pm PST
Subject: Wordperfect formatting

[From Dag Forssell (921119.1100)] len lansky (921119.0816)

>Thanks also to Bill Powers re idea about margins. I am
>working on it; the problem is getting rid of hard returns--
>I think I or our system is screwing up going back and forth
>between ascii and wordperfect.

I work in WordPerfect. Bet your problem is WP. What is your default printer? (Shift F7)
What is the default font? (edit) Pick a small one! 12cpi or 10 points. What are default margins? See setup (Shift F1) initial settings (4) and initial codes (5).

When you call up an ascii file - List Files (F5) and select it (Ctrl F5) (best) or (1) your document gets formatted as specified in all these places and hard returns inserted as necessary to conform.

There is an alternate way to call up an ascii file, which may solve the immediate problem. Directly from the screen, (Ctrl F5) Dos text (1) Retrieve CR/LF to [Srt] (3)
Document: C:\etc.

Remember, you can install any printer into the program, if you want to see what another printer selection can offer you.

The above remarks apply to WP5.1 and 5.0.

Best, Dag

Date: Thu Nov 19, 1992 4:28 pm PST
Subject: & action)

(penni sibun 921119.1300)

We obviously don't need to figure out how words are recognized or produced; we can just say that they are. The associative part looks easy to implement in a limited situation.

I think we could make up a working model that could do the following: given the name of an object, turn and look at it; given a picture of an object, name it and turn to look at its prototype in the room; given a direction in which or a location at which to look, know that the object seen is a familiar object and name it. We would have to stipulate object recognition. Perhaps we could work in kinesthetic body configuration sensing, but we would have to stipulate, for the time being, conversion of visual images to objective space. But even with these limitations, I think this might be an interesting model. I invite anyone who is interested to be the first one to get such a model working.

chapman's system Sonja does just this sort of thing. chapman also finesses the hard problems you suggest finessing, viz., object recognition and language use.

cheers. --penni

Date: Thu Nov 19, 1992 5:07 pm PST
Subject: what's in a code?

(penni sibun 921119.1400)

[From: Bruce Nevin (Wed 921118 12:07:22)]

My response was that in the party patter in fact considerable linguistic information is in the language used, but that it is not information that the participants care much about;

sure, there's ling info in the lang, but the users didn't put it there. it's (often) an artifact. i claim this is frequently true when we speak our native language; it's quite obviously true in a foreign lang, as illustrated by this excerpt from a recent article about japan by dave barry:

When we pulled up in front, three women in kimonos came out and began bowing and saying things in Japanese and picking up our luggage. Using our Japanese skills, we said "thank you" or possibly "good night," and we bowed, and they bowed some more, which was not easy for them to do while holding our luggage.

Our maid, who was wearing a kimono and a beeper, came into our room about 30 seconds after we arrived, speaking in the very high, sing-song voice that Japanese women often use when they're speaking to somebody in authority. "Hai domo!" she said. She said "Hai domo!" to us a lot. As far as I was able to determine, "Hai domo!" means "Yes very!" We came to think of her as the Very Lady.

and in the flirtation little information is in fact transmitted, however much is imagined, and that the most important thing communicated is interest in and availability for negotiating mutual agreements that have a more explicit foundation.

it is clear that we will never find ourselves flirting, since we have radically different ideas of what's involved.

>in my description of flirting, i gave specific instances of
>information being conveyed (eg, ``i like it when you do that'').

You are imagining that you have conveyed the information that is in the sentence "I like it when you do that." He may very well be imagining something else, such as "she thought that was funny".

i don't go around imagining that i convey info in sentences. i gave you a verbal gloss of a nonverbal communication. of course confusions such as those you suggest occur, and either get subsequently resolved, perpetrate tragicomic consequences, or are irrelevant. these exact same issue arise *all the time* in verbal communication too.

probably nothing even as specific as your really rather vague "I like it when you do that"

that's not vague at all--it's terribly concrete. it looks vague when we write it down, cause there's no context, and every word in it has deictic/indexical qualities. in context, the whole expression refers to--the context, which is right in the participants' face(s). maybe it takes some negotiation to figure it out precisely, but the space of possibilities is pretty small, being the-here-and-now.

Even the more specific "I liked it when you just did that" has this problem.

this is more specific cause it's more inflected??

Conversely, this is why language is good for error-free transmission of linguistic information,

what is ``error-free transmission of linguistic information''?

but not particularly good for communication, in the sense of communication that happens in flirtation and is restrained in party patter.

--and if it's unrelated to the bulk of communication, then why should we focus on it?

cheers. --penni

Date: Thu Nov 19, 1992 5:09 pm PST
Subject: Canon ahoy

[From Rick Marken (921119.1300)] Bill Powers (921119.1000) --

>Rick Marken (921118.1300) --
>It is important for a loose canon not to fire when pointing inboard.

Ay. And I'd like to apologize to the mates for me little fit, thar. You'd think I was ridin' HMS StimulS-R instead of the SS PCT.

Best reagrds Rick

Date: Thu Nov 19, 1992 6:04 pm PST
Subject: images, etc.

[Avery Andrews 921129.0930]

So the way it looks to me at this moment, the perceptual reference `signals' that you would derive from the text to tell you where to go would actually have a lot in common with the `mental models' postulated by people like Johnson-Laird. Their supposed properties are (to my mind at least) a rather unclear combination of imagistic & and logical properties, & which Bill's view of things might help to clarify.

So next question: how do we distinguish between:

a picture of a cat chasing a dog

a picture of a dog chasing a cat

Maybe bonkers idea. What if there were two areas, an event area where patterns of neural activity corresponding to verbs occurred, and an entity, where patterns corresponding to the entities occurred, with the normal situation being that the event pattern (thought of as a static configuration of some fibers firing and others not) coexists with a sequence of entity patterns, in a standard order (presumably Agent - Patient, though there's lots of other roles to fit in).

These would get correlated with perception via complex events being scanned in a characteristic order. E.g., `chase' refers to a perception whereby on the entity level, you look first at the one behind and register its features on the entity level, and then do so for the chasee, while `flee' scans in the reverse order.

Remembering a complex perception would then amount to remembering a sequence of scanning activities & attendant results. How this is to be achieved I don't have any clear ideas. But, for what it's worth, the inner representations I'm suggesting have pretty much the same organization as appears overtly in manual sign language, where signs representing actions are articulated simultaneously with those designating participants (coded as pointing gestures to positions in space, which have been previously been associated with specific referents).

Avery.Andrews@anu.edu.au

Date: Thu Nov 19, 1992 6:25 pm PST
Subject: one-neuron ECSS

This is a test of an automated "send" script, but I thought I'd make it useful while I'm at it.

Martin Taylor (921119.1100) --

>If it is true, as we discussed some months ago, that an ECS can be
>represented by a single neuron, then that neuron seems to do
>exactly the combination your are talking about. Each dendrite

>serves the reference-perceptual input-comparator functions, but
>they combine into a single output function.

An ECS can be represented by a single neuron, but I don't know of any examples where it is. The simplest control systems I know about, the spinal reflexes, involve one neuron for a sensor, one neuron for a comparator (the spinal motor neuron) and a muscle for an effector. Furthermore, as in the interchange with Oded today, these simple control systems operate in parallel bundles so the effective control system probably involves dozens to hundreds of neurons, plus the muscle fibers.

I would expect that a control system that senses and controls a configuration requires dozens to hundreds+ of neurons even to carry out a simple pattern-forming input function. The comparator can be simple, but the output function may also need a rather complex computation needing more than you can do with one neuron.

On the other hand, a neuron with 5000 input synapses might be able to accomplish some pretty amazing analog computations. And that's just the average number in the brain (roughly).

Best, Bill P.

Date: Thu Nov 19, 1992 6:46 pm PST
Subject: association & g-ma neurons

[Avery Andrews 921120.1145]

This may be pretty ignorant, but here goes anyway. It seems to me that a characteristic feature of classical 'connectionist' systems is precisely that they can in a sense recognize configurations without having a 'grandmother neuron' for each config. that is recognized. E.g. you can have a bunch of input & output wires, with an arbitrary association function connecting input to output patterns:


```
---  
--->| |--->  
--->| |--->  
--->| |--->  
---
```

For example on the input lines might be visual/auditory/olfactory properties, on the output lines kinaesthetic reference levels, such that for input combinations typically associated with the presence of leopards, the output levels are those associated with running away, etc.

This sort of thing should also be possible for connecting verbal & nonverbal perceptions. To account for actual learning of words, it would have to be 'programmable' very quickly (on the basis of a single experience), but I think there might be a preadaptation to this sort of things.

Imagine a very simple critter whose sense inputs are like e-coli's (e.g. a vector), but which also has a 'pain' perceptual system that registers damage. We imagine the sense vectors and the pain wire going into a black box whose output will be labelled 'fear':

```
|  
| pain
```



The idea is that whenever the pain wire fires, something happens inside the black box such that, subsequently, the closer the sense-vector gets to what it was when the pain wire fired, the larger a signal comes out along the fear wire. This signal then is then handled in such a way is to typically induce behavior that will alter the sense-vector. (If the organisms motor system is e-coli style, the fear signal would contribute additively to the tumbling rate).

Stick a lot of these things together in the right way, & you might be able to get a fast associative learner with no grandmother neurons.

& of course, for the optimists, a similar picture could be drawn with wires labelled 'pleasure', etc., where the output wire tends to do things to cause the sense-vector to get closer to what it was when the pleasure wire fired (e.g. reduce tumble-rate).

Avery.Andrews@anu.edu.au

Date: Thu Nov 19, 1992 9:44 pm PST
 Subject: Promoting PCT, Version 13

[From Dag Forssell (921119.2100)] Ray Jackson (921118)

Re: WHY A THEORY-BASED LEADERSHIP PROGRAM

>I hope you find these comments encouraging and useful.

I do indeed. Praise in general does not help. Specific suggestions do. You gave me many here. I have used several.

I have some trouble with the 2-dimensional metaphor. It has an engineering meaning. I have taken your intent, but not the metaphor. See below.

Your rewrite is 50% Mary, 50% original. I have taken a careful look by highlighting Mary's in yellow and yours in blue. You are a good editor, writer and you demonstrate understanding. Well done.

>In its elaborated form, a Hierarchy of Perceptual Control Theory ...

This seems wrong, however. You don't have a hierarchy of theory. Back to Mary or strike the hierarchy part. Why confuse the issue. Talk PCT only.

What quotes from Kuhn do you have in mind? What do you present?

Below is version 13 (lucky number?) of solicitation letter for CEO. This is the first version to directly address some likely concerns up front. Comments appreciated as always. Note: Only two pages. No enclosures.

PURPOSEFUL LEADERSHIP
(Logo)

Dag Forssell
Address etc.

William T. Powers, Salesman
PCT Manipulators, Inc.
73 Ridge Road CR 510
Durango, Co 81301

November 19, 1992

Dear Mr. Powers:

The recent article: ALLIED-SIGNAL'S TURNAROUND BLITZ (Fortune, Nov 30) recounts a success story of corporate revolution and turnaround. It spells out several issues important to the revolution. I have grouped them in three categories:

1) QUALITY/PRODUCTIVITY:

Issues: Total Quality immersion. Quality should be routine; no bureaucracy.

Comment: The absence of a clear overall framework of understanding for TQ makes it hard to teach. "Perceptual Control Theory" (PCT) offers a framework for understanding TQ, and for the people issues of TQ. Quality is one dimension of perceptual control and must be a routine result. It cannot be "controlled" by any bureaucracy.

2) PEOPLE:

Issues: Human Resources top to bottom review. Culture change. Depth of talent available. Employee wants. Curriculum. Training with real problems.

Comment: The absence of a clear and valid theory of psychology in our society makes people issues complex for all of us. PCT offers a psychology based on clear, detailed and tested engineering concepts. Theory and practice go hand in hand. This theory sheds new light on focus group results, training curriculum and leadership. It explains the essence of empowerment.

3) SPEED:

Comment: An in depth understanding of PCT gives you the capability to evaluate and consolidate training programs, reduce confusion and save time.

PCT offers a new explanation of how the mind works; why people do what they do. When you and your associates understand PCT and decide that it makes sense, you will look at yourself and others in a different way. You will know what can and cannot be done in harmony with others. This understanding and our methodology leads to cooperation. Productivity, quality, relationships and satisfaction will all improve.

PCT can explain a wide variety of phenomena of everyday experience because it goes beyond the predominant focus on cause and effect to explain the consistency of outcomes through varying means.

William T. Powers

November 19, 1992

Page 2

Understanding people no longer has to be complex and confusing, a matter of luck or a gift. PCT can be taught in simple form with a comprehensive management application in one day and in more detail with leadership applications in three.

The new perspective does not invalidate any wise common sense observation or practice. It does provide an enhanced understanding of seemingly intractable problems. It gives you new diagnostic tools and shows why cookbook rules for behavior (programs which tell you what to do under certain circumstances) do not always work.

We provide step-by-step explanations of all aspects of Perceptual Control Theory with complete, clear illustrations and many participative demonstrations.

Participants learn a universal, simple and easy to remember methodology we call Teaching Effectiveness, based on the theory in an obvious way, that helps them think and deal effectively with real people problems as they occur.

Applications that follow directly from this approach are: conflict resolution; performance reviews; team development and non-manipulative selling. Role play practice with a real problem and class discussion of experiences are provided for each of these.

Applications where the theory provides insight are: leadership understanding; goal structuring; vision & mission statements; a framework for total quality management.

A participant commented:

We have had quality management training. We were taught to focus on the problem rather than the person. You have plugged a big hole for me by showing me how to deal with people problems rather than avoid them.

People trained in the technical sciences appreciate the scientific approach and elegant simplicity of the theory, and everyone is able to begin applying the principles as soon as they understand the model and have had some instruction and practice.

A personal meeting to demonstrate the concept and explain the theory would be best, but I will be pleased to send you your personal selection of introductory material. You may request an audio tape with script to explain the nature, background and content of our programs. It includes an active do-it-yourself demonstration of the basic concept. The demonstration allows you to find out if your associates can recognize simple control in action. (I bet they can't). You may also choose from a selection of papers. This material will allow you to evaluate the validity of this information and its potential impact on PCT Manipulators, Inc. Just give me a call.

You may find the demonstration enlightening and entertaining. This win/win program will increase the understanding and effectiveness of anyone who deals with people.

Mr. Powers, if you are striving for a corporate revolution, my revolution can support and accelerate yours. I would like an opportunity to show you how.

Sincerely,

Dag Forssell

The condescending part of my draft, deleted yesterday:

A revolutionary CEO:

Strives for rapid improvement.
Accepts personal responsibility.
Voices personal convictions.
Questions conventional wisdom and
breaks with it where it is wrong.

Keep them comments coming - on the net.

Control well, Dag

Date: Thu Nov 19, 1992 10:20 pm PST
Subject: your page today
Ray - Direct
Hello Dag,

I wanted to send you a note to let you know I got your page today; however, today I taught one of my classes until 3:30, then I drove across Phoenix to the classes I take at Grand Canyon University, where I spent from 4:00 until 10:00 becoming an educated individual.

Anyway, it's about 11:30 here, which it seems is often the only time I can do anything of any consequence. I'll try to call you tomorrow (Friday), or at least over the weekend, okay?

As I said in the post..."keep at it",

Take Care Ray

Date: Thu Nov 19, 1992 11:46 pm PST
Subject: Re: Promoting PCT

[From Ray Jackson (921119.2345)] for Dag Forssell (921119.2100)

Re: WHY A THEORY-BASED LEADERSHIP PROGRAM?

My, you're up late. I'm always thrilled to help your efforts. Just a few minor comments and suggestions (nothing major, I'll leave those to Bill).

>I have some trouble with the 2-dimensional metaphor.

Understandable; I find that metaphor more useful in the education arena, when I can explain it standing up, using my hands, and looking in people's eyes. I like what you did to simplify it later in your letter.

>In its elaborated form, a Hierarchy of Perceptual Control Theory...

I agree, PCT only for initial introductions in the business world; in fact, I only put this in the summary as an introduction for the graduate class so they understood what HPCT was. I did double-check -- I had taken it out of the version I distribute at the plant.

>What quotes from Kuhn do you have in mind? What do you present?

Mostly those where he summarizes his findings. Often, I read aloud entire paragraphs from his book. I'll put those citations out in my next post. In that context, I also often use Joel Barker (The Power of Paradigms, book & video) as an introduction to a discussion of mindsets. Although Barker is shallow and inconclusive, he's thought provoking when followed with valid material.

>Note: Only two pages. No enclosures.

Beautiful. It can be read while another meeting is going on.

>The recent article: ALLIED-SIGNAL'S TURNAROUND BLITZ (Fortune,11/30)
>recounts a success story of corporate revolution and turnaround. It
>spells out several issues important to the revolution. I have
>grouped them in three categories...

This is great; talk about working with THEIR error signals!

>PCT can explain a wide variety of phenomena of everyday experience
>because it goes beyond the predominant focus on cause and effect to
>explain the consistency of outcomes through varying means...

>Understanding people no longer has to be complex and confusing, a
>matter of luck or a gift...

This is EXACTLY what managers with the appropriate error signals wonder while they're showering or having their morning coffee.

Good work, Dag.

Let me know how I can help (on video, my best side is my right, in case you need help when CareerTrack or 60 minutes calls).

Best Regards, Ray

Ray L. Jackson 602-963-6474
3613 W. Saragosa St. Chandler, Az 85226 attmail.com!rljackson

Date: Fri Nov 20, 1992 9:42 am PST

Subject: S-R theory and Neural nets

[From Bill Powers (921120.0830)]

Avery Andrews (921120.1145) --

You're reinventing S-R theory. Your proposal is essentially how neural network people are trying to use neural nets -- to connect inputs to behavioral outputs. The assumption behind this approach is that for each discriminable situation, there is an appropriate output that will take care of it. This assumption is false, and is the basic reason for the existence of PCT.

>The idea is that whenever the pain wire fires, something happens
>inside the black box such that, subsequently, the closer the sense-
>vector gets to what it was when the pain wire fired, the larger a
>signal comes out along the fear wire. This signal then is then
>handled in such a way as to typically induce behavior that will
>alter the sense-vector.

The catch in this is "in such a way as to typically induce behavior that will alter the sense-vector" [in the right direction and by the right amount]. The whole problem is how to induce just that qualitative and quantitative action which, when added to all the other influences that are also acting on the sense-vector, will move the sense-vector toward a particular state from any initial state.

It's true that what is needed is to move the sense vector in the right direction. But that is an OUTCOME, not a MEANS. The means employed must be variable, and generally is of a different nature from the outcome. Suppose the desired outcome is to move "away" from a threatening (fear-inducing) object. The means of doing this is to alter the way the legs are thrusting against the ground. But doing that in the right way requires

(1) knowing the orientation of the body (do you just run or do you have to turn around first, or run backward?),

(2) knowing the composition of the footing (are you on slippery mud, in sand, or standing on concrete, or on an upslope, a sideslope, or a downslope?).

(3) knowing the proximity and location of obstacles (is there a tree in the way? Will I run off of a cliff?),

(4) knowing the current velocity and direction of locomotion (am I already running in the right direction, or at right angles, or toward the threatening object?),

(5) knowing the present bodily configuration (am I lying down, sitting down, kneeling, standing? Is my head pointed in the same direction as my body? Is the object centered in vision or in the periphery?),

(6) Knowing the probable reaction of the object to my running away (will that encourage it to attack? Can I run faster than it can? Would I do better to run at it and scare it away?),

(7) Knowing my relative strength and weaponry (Should I just kick it aside? Shoot it? Grab it and push it away?),

and so on.

Under the S-R concept, it's necessary for each possible factor in the environment that might bear on the choice of appropriate output to be sensed and taken into account in the computation of output. It's also necessary that every factor capable of altering the sensory outcome be known -- that is, that the causes of possible disturbances be sensed. But the causes of many important disturbances are insensible -- the first we know of them is that the sense vector in question starts to change, apparently spontaneously. Such a spontaneous change implies that the current action is incorrect -- but there is no previous cue to tell the computer what action WOULD be correct.

So an S-R model that would actually account for behavior in the way you suggest would be EXTREMELY complex. It would have to have precisely calibrated senses which would maintain their accuracy over long periods of time. It would need to infer things it could not sense. And it would have to compute the inverse dynamics of all physical output effects, taking into account (by sensing them) the masses of the moving parts, the influences of external forces like gravity and inertial forces, the friction in the joints, the state of metabolic energy supplies near the muscles, and the interactions among all muscle forces being used at the same time. It would have to be acquainted with all the relevant laws of physics and chemistry and geometry.

A control system, on the other hand, needs to sense only the outcome that is to be controlled. It must be connected so that for any change in the monitored outcome, it can exert an output effect directly on the outcome, opposing the change. There is no need to compute inverse dynamics. There is no need to sense the causes of any disturbance of the input. There is no need to predict the behavior of such causes. There is no need to understand physical laws. A control system is, by a wide margin, the simplest system for accomplishing the things we observe to happen. It may still be complex enough to puzzle us, but the problems we face in modeling behavior with control systems are NOTHING compared to what we would face in building a working model of an S-R system -- that had to work in a real environment.

Grandmother neurons are not a problem. After all, we already know about "line tilted 22 degrees from the vertical" neurons, and "this is a particular face" neurons (in sheep, yet). Of course beneath those neurons where we measure signals corresponding to such entities there is a lot of computation -- the neuron is just an output neuron. But there's nothing in principle or in experimental fact to rule out the idea that we might measure a signal at a particular neuron when and only when grandmother is present. The "grandmother neuron" objection is one of those facile concepts that is famous because it's so well known.

The disk will go in the mail today.

Wolfgang Zocher --

I have tried to reply directly to the return address on your posts and I get "illegal host address."

I am VERY interested in your program for assembling control systems using modules. If you will send me the source code I'll see if I can get it running, and offer any suggestions that seem useful. If we can work out something with the required properties, I'll switch to using it in my Primer series.

Bill P.

Date: Fri Nov 20, 1992 10:34 am PST
Subject: Grandmother cells

[From Rick Marken (921120.0800)]

In a personal reply to my perceptibly rude public reply to him, Martin Taylor did answer my question about the "Grandmother cell" -- the question being "what's wrong with it"? I will take the liberty of posting Martin's very helpful description of the Grandmother cell problem along with my apologies for responding testily to his previous post. Sorry Martin -- really.

Now, on to Grandma.

Martin Taylor (in a private post to me -- 921118) says:

>The argument that has always been used, I think successfully, is that
>there would have to be too many of them to fit into a brain. That
>problem is finessed by distributed representations, since the amount
>that can be stored is not limited by the number of storage places, but
>by the differences that can be discriminated among vectors.

I think this is an important point because it reveals another area (perception modelling) which has been strongly influenced (tainted?) by the digital computer metaphor of the brain. Your statement of the "Grandmother cell" problem implies that perceptions are either present or not; the grandmother perception is present if the grandmother cell fires (above a threshold level, perhaps) and is absent if it does not (or fires below threshold). This implies that perceptions are either/or phenomena -- and they are indicated by the presence or absence of "signal" in a cell. This model does seem to require a different cell for every different perception; and this could, presumably be a huge number of cells.

The "Grandmother cell" is not a problem for the PCT model because perceptions are not either/or phenomena; they are VARIABLES (just like the rate of firing in neurons). So you don't need a different cell for every different, say, configuration you can experience; just cells that respond in proportion to the degree that certain configurations are present; a roundness detector cell, for example, where the firing rate is low for something like a compressed ellipse, increases as the ellipse becomes less eccentric and peaks out when it is a perfect circle. Perceptions vary continuously -- and there are many different types of perceptions that vary within type (so to speak); warmth (a sensation) varies as you adjust the water to the right temp; grip (a configuration of force sensations) varies as you adjust your hand to the steering wheel; transition varies smoothly and can be changed easily by varying the pressure on the gas pedal. All these different sensations, configurations and transitions can be represented as different rates of firing in individual neurons. One can even perceive different degrees of "grandmotherliness" (probably a category level perception) running from Joanne Woodward (least like) to Barbara Bush (most like).

At some level in the hierarchy, perceptual signals do seem to be treated in either/or form; it might even be that there are perceptual variables treated in either/or form at all levels of the hierarchy; in these cases, each neuron does function like a grandmother cell; but there is no reason to think that the number of such "categorical perceptions" is anywhere near the number of perceptual cells available at each level of the brain. A few hundred phomemes, a few hundred thousand categories, etc. With something on the order

of 10 billion neurons (and maybe 5 billion available for perceptual signals) I don't think there is a real storage problem.

Anyway, the presumed solution to the Grandmother cell storage problem -- distributed representation -- is only a solution in the minds of the people watching these distributed systems; at some point, some resource (a cell or cell assembly) is going to have to represent the state of the distributed representation to someone other than the person watching the model -- ie. there will have to be a perceptual signal that represents the state of the distributed representation -- and that requires storage.

Best regards Rick

Date: Fri Nov 20, 1992 10:43 am PST
Subject: Re: association & g-ma neurons

[Martin Taylor 921120 11:30] (Avery Andrews 921120.1145)

>The idea is that whenever the pain wire fires, something happens inside
>the black box such that, subsequently, the closer the sense-vector gets
>to what it was when the pain wire fired, the larger a signal comes out
>along the fear wire. This signal then is then handled in such a way
>is to typically induce behavior that will alter the sense-vector. (If
>the organisms motor system is e-coli style, the fear signal would
>contribute additively to the tumbling rate).

>
>Stick a lot of these things together in the right way, & you might be
>able to get a fast associative learner with no grandmother neurons.

This looks to me so much like Bill's reorganization of the perceptual input function that I can't tell the difference. I had originally thought that Bill was applying reorganization only to the output links among ECSs, but some time ago he mentioned experiments he had been carrying out on e-coli type modification of the perceptual input function as well. And since "pain" would not be an input to the ECS that is reorganizing, but might be a kind of perceptual consequence of an error in a critical (intrinsic) variable, the line you have labelled "pain" would actually be a reorganization signal that occurred because of the intrinsic error, rather than the pain perception itself.

But yes, there would be lots of those things, and, I claim, no need for a "grandmother cell" to control all the many perceptual attributes of grandmother. The only necessity in the hierarchic control system is that there be ECSs that control perceptual signals whose control happens to be effective in maintaining critical variables near their reference levels. The more complex the living control system, the more adaptable it is, by virtue (speculation mode) of being able to control perceptual signals that correspond to more complex CEVs in the real world.

(Bill Powers--whose send script doesn't yet include a date stamp--921119 1640)

>An ECS can be represented by a single neuron, but I don't know of
>any examples where it is. The simplest control systems I know
>about, the spinal reflexes, involve one neuron for a sensor, one
>neuron for a comparator (the spinal motor neuron) and a muscle for
>an effector.

There's a critical point about this particular identifiable control system, and that is that its peripheral components are not ECSs. What I mean is that the sensor neuron is essentially different from the final neuron that activates a mechanical muscle, and therefore the lowest level ECS that connects peripheral sensors with muscular actions MUST have at least two neurons. As you say, it actually seems to have three. But when you get deeper into the system, the outputs from any ECS could come from the same neuron that receives the inputs to it, meaning that the whole ECS could be in one neuron. Since we do not (so far as I am aware) know of any physiologically identified higher-level ECSs, I don't see where there is evidence one way or the other that the smallest unit that acts as an ECS is smaller or larger than a neuron. But the smaller the unit, the weaker the numbers argument against grandmother "cells" (really ECSs).

>Furthermore, as in the interchange with Oded today,
>these simple control systems operate in parallel bundles so the
>effective control system probably involves dozens to hundreds of
>neurons, plus the muscle fibers.

This I do agree with. It's part of the foundation of my arguments about what we might call "subsymbolic PCT" if the word "symbol" didn't have an unfortunately double meaning as a percept and as an analytic concept in PCT. These dozens to hundreds of overlapping outputs (and inputs) are unlikely to form strictly isolated modular groups, but are more likely to be interwoven and to blend into each other. Just as control of a low-level percept is effected by a separated muscular output, and as control of intrinsic variables (in your model) is effected by separated control of apparently irrelevant perceptual signals, so may the "core" percepts in a modular bundle be partly controlled by actions of peripheral members of the bundle.

Martin

Date: Fri Nov 20, 1992 1:53 pm PST
Subject: Re: Grandmother cells

[Martin Taylor 921120 13:40] Rick Marken 921120.0800

Apology accepted, Rick, and no (permanent) offense taken.

>Anyway, the presumed solution to the Grandmother cell storage
>problem -- distributed representation -- is only a solution
>in the minds of the people watching these distributed systems;
>at some point, some resource (a cell or cell assembly) is going
>to have to represent the state of the distributed representation
>to someone other than the person watching the model -- ie. there
>will have to be a perceptual signal that represents the state
>of the distributed representation -- and that requires storage.

This is at the heart of the disagreement between the AI people and the connectionists, and it is interesting to see it rearing its ugly head within PCT. The question in my mind is: why should it be true that

>at some point, some resource (a cell or cell assembly) is going
>to have to represent the state of the distributed representation
>to someone other than the person watching the model -- ie. there
>will have to be a perceptual signal that represents the state
>of the distributed representation...

Why does it need to be represented, within the PCT view of the world? It seems to me that all that need happen is that the various perceptual (scalar) signals within the complex be controlled. The distributed representation is a representation of "something" only if there is someone outside that needs to represent it and (presumably) label it, perhaps not verbally. That includes the person whose brain includes the complex, seeing themselves as if from an outside observer. For acting in such a way as to control the complex, a single higher-level ECS will not do the job, because the complex has more than one degree of freedom. You need many higher-level ECSs, and then they themselves form a complex.

Sure, you can have "grandmother" ECSs, and as Bill pointed out, specific face detectors have been found in sheep. I never argued that you couldn't, just that for many purposes you needn't. That's what I meant when I originally said Bill's point was unnecessary. (Last night, it occurred to me that perhaps you thought I meant it was unnecessary for Bill to have made the point. That would have been discourteous and would have justified the tone of your original response; I meant it was technically unnecessary.)

I was thinking a bit about conscious perception. Whatever it may be and wherever it arises within the hierarchy, it must relate in some way to perceptual signals. But it is subjectively very highly multidimensional, and I don't see how this kind of subjective multidimensionality can be reconstituted from one scalar signal, no matter how complicated the function from which that one is derived. That leads me to think that at least the conscious aspect of perception must be derived from a complex of scalar perceptual signals. If so, then why should the set of perceptual signals relating to one perceptual "object" such as "grandmother" be required to be combined into one scalar signal, at any point at all in the process? Again, it seems necessary to combine them only if you are going to "represent" the object as something (a symbol, a word, a non-verbal label).

Scalar perceptual signals may well be a fundamental core for perceptual control, but they don't happen one at a time. I guess that's the bottom line for this argument.

Martin

Date: Fri Nov 20, 1992 2:00 pm PST
Subject: Perceptual Complexities

[Oded Maler 921120]

(Bill, Martin and Rick (the true believer) 921119)

The question whether some vector in time-space can be considered as an aggregate scalar, depends the granularity (again, spatio-temporal) of the model. In Bill's example of the stretch reflex, if you look from the point of view of single muscle cell or a neuron, there is no such a thing as "tension along the muscle group", it is only in the head of the beholder, there is no single probe than you can put inside that will give you the value of this variable. Bill claims that when the whole is indeed the (statistical) sum of its parts (= positive linear combination) then the virtual signal can be considered as "really" existing in the model.

* * *

The question whether you can servo two low-level signals that are combined thru action, in such a way that it looks as if you control their combination (virtual signal), is, I think, dual to the question being able to control the real signal which is their combination by changing their reference signals. I'm not in the mood to try to prove it, but I'd appreciate counter-examples.

* * *

The difference between real and virtual signals (to the extent it really exists) might have very interesting epistemological consequences. If you adopt, say, voltages along axons as your "atomic" units of perceptual variables, you might distinguish between world views of different persons according to whether some CEV has or has not a dedicated correlated real signal. It is also might help in understanding the role of words, not as designators of meaning, but as a set of abstract common entities the each of the speakers of the same language has a dedicated signal for most of them. This way, words strengthen the illusion of common external world, independent of the different low-level hierarchical structure of each individual.

Bon weekend a tous

--Oded

Date: Fri Nov 20, 1992 4:31 pm PST
Subject: Movin'

I'm moving to a new office and a new e-mail address. So I will be unsubscribing from csg-1 briefly and , with any luck (for me; maybe not for all of you), I will be signing back on soon with my new address (in my sig, below; also, note, new work phone #).

Best regards Rick (the true believer .. er .. modeller) Marken

Date: Fri Nov 20, 1992 4:39 pm PST
Subject: Complex perceptions

[From Bill Powers (921120.1300)] Martin Taylor (921120.1340)]

You replying to Rick:

>The question in my mind is: why should it be true that
>>at some point, some resource (a cell or cell assembly) is going
>>to have to represent the state of the distributed representation
>>to someone other than the person watching the model -- ie. there
>>will have to be a perceptual signal that represents the state
>>of the distributed representation...

The "someone else" is a higher level system. This discussion is making me go back through some old musty trains of thought that eventually led to HPCT.

When you have a set of ECSs at a given level, each one provides a controllable perceptual signal (there are also uncontrolled ones, in general, at the same level). A copy of the controlled signal is available to higher-level systems. The question before us is whether

these independent control systems can produce a "complex" of perceptual signals that has some specific pattern in it, so the whole ensemble of perceptual signals has some significance other than that of each individual signal.

The basic answer of HPCT is no, that isn't how hierarchical perception is supposed to work in this theory. To see why not, consider controlling for some linear combination of the perceptual signals by setting each signal to conform to one element of a pattern. To do this it would be necessary for a higher system to set the reference signals for each of the lower ECSs to satisfy the same linear combination. In the event that the linear combination (or nonlinear for that matter) were made up partly of the uncontrolled perceptions, there would be no way to specify reference signals that would satisfy the specification.

In the hierarchical model, a linear combination of the lower level signals would first have to be sensed and converted to a single signal by the higher system's perceptual function. That function would contain exactly the linear combination to be controlled -- that is, the weightings would represent the coefficients and the computing operations like adding, subtracting, multiplying, etc., would do the combining. This single signal would be compared against a single reference signal, and the single error signal would fan out to the lower reference inputs with appropriate positive and negative weightings. But these output weightings now do not represent the linear combination itself, but only signs of effects. As long as the sign is correct in each output path for making the higher-level error smaller, no output weighting is strictly necessary (although it can help).

The important thing to see here is that there is no one distribution of the lower-level perceptual signals -- no one "pattern" -- that is necessary for achieving a particular value of the higher-level scalar perceptual signal. What the higher system is controlling is not a particular set of values of the lower-level perceptions, but a particular multidimensional relationship among them, or more precisely a function of them. There are many ways in which those lower-level signals can vary without altering the value of the function and hence the higher-level perceptual signal. Changes in the pattern of lower-level signals that do not alter the higher-level perceptual signal are simply not controlled (except by individual systems at the lower level). It is not necessary for all the perceptual signals at the lower level to be under control, nor is it necessary to protect them against variation. If uncontrolled perceptions vary, the controllable perceptions can generally be altered to restore the value of the function to its specified level.

In connecting this picture to conscious experience, confusion arises because we can be aware of more than one level of perceptual signals at once. When I look in the basket hanging at the top of the stairs, I see a letter in it. Ah. A letter in the mail basket. Then some higher-level perceptual functions kick in, and now I see, oh crap, an UNMAILED letter that was supposed to go out today, and Mary has already left for town (don't worry, Mary, I'm making this up).

Now where was that "unmailed-ness" when I first looked at the letter? Objectively, I was looking at an unmailed letter from the very start. But I didn't perceive its unmailedness. The mere existence of the configuration and relationship perceptions was not enough. Some higher-level system had to perceive these elements, put them together into a perception of the state of mailed-ness, compare that with the desired state, and come up with an error big enough to attract my attention. THEN I experienced the significance of the pattern of lower-level perceptions. I still perceived the configurations and relationships. But now I also perceived a new signal. If I had intended to delay mailing it, I might then say good, she remembered not to take it.

This is why I say that it's not enough for a pattern to be implicit in a set of variables. An implicit pattern can have no meaning to a higher system. In fact, thinking of patterns in collections of perceptual signals is backward to the HPCT way of modeling perception. I know it's customary in conventional approaches to think of the pattern than is present and try to design a recognizer for it. But in HPCT there is no "inherent" pattern.

The HPCT view is that the pattern is imposed from above. When a higher perceptual function becomes organized, it receives signals that are lower-level perceptions, puts them through some sort of computing function, and delivers a perceptual signal carrying the scalar value of that function. The form of the computing function determines completely how the perceptual signal will vary as the lower-level perceptions individually vary. EVERY function will produce a perceptual signal of some magnitude. EVERY function with multiple arguments defines some sort of invariant. ANY ARBITRARY function of the lower-level signals can be brought under control, provided only that the value of the function can be systematically affected by altering reference signals going to lower-level systems.

No matter what the form of the function, it defines a pattern in the lower-level signals, an invariant. Other higher-level systems can employ different forms of perceptual functions using the same lower signals as inputs or arguments; they will perceive and control a different pattern in those same lower-level signals, at the same time (if there is sufficient orthogonality).

We do not learn to perceive patterns that are "really there." We learn to perceive patterns that are useful to control. Somewhere, way off in the philosophical distance, there may be a deep principle involved, to the effect that the patterns that prove to be the most useful to control, in terms of the effect of doing so on critical variables, happen to relate in some way to something that is in some sense "really there." But organisms do not know what is really there. They know, or learn, only that seeing the world in particular ways, and controlling what is seen relative to particular reference levels, feels better than not doing so. No philosophy, except in some human beings who have taken up the habit.

I think that the idea of complex perceptions, vectors that have meaning in themselves, is a vestige of an older way of modeling behavior based on a realist epistemology. We certainly experience multiplicity of perceptions. But until we form the right perceptual functions at a higher level, we do not perceive pattern in them, and none of the implicit patterns, infinite in number, has any effect.

An exercise that I went through frequently in developing definitions of levels of perception was to try to imagine what the world would be like if a given level were the top one. Suppose, for example that the top level was configuration (as it must be in some sufficiently simple organisms). Watching someone blow up a balloon, one would see "Limp rag, limp rag, balloon, balloon, balloon, (POW!), no balloon." There would be no perception of transition, no sense of the balloon getting larger and larger. There would be no relationship between one size of the balloon and another size of it. One state of the balloon would have no causal relationship to any other state. There would be no logical relationship between the size of the balloon and its bursting; in fact its bursting would not constitute an event. There would be only shapes.

By using this ploy with various levels of perception, you can begin to see what a higher-level perception is. It's a new signal, appearing in experience along with all the other lower-level signals that are already present. When you see the balloon increasing

in size, you still see the shape at every moment and the color and the edges and curvatures. But a new signal appears that is not present when the balloon remains at constant size. This new signal is the sense of transition, of growing. It has no other character. It is simply a signal that occurs when someone blows more air into the balloon. You can't take this signal apart into elements; it is unitary. This signal IS what we mean by the term "expanding." It does not refer to anything else; it is that to which we refer with words denoting a change in size. We refer to a small magnitude of this signal by using words like "slowly" or "gradually." A largemagnitude is what we mean by "rapidly." Zero magnitude is what we mean by "constant" -- a word that has no meaning unless you're capable of perceiving transitions.

What you end up with is a picture of your own world as a collection of attributes or aspects each of which is just a single variable amount of something. This is true even for such an abstract perception as honesty. You can enumerate lower-level perceptions from which this sense of honesty is derived, but the honesty-signal itself is just a signal, with no fine structure.

One perception, one signal.

Best, Bill P.

Date: Fri Nov 20, 1992 5:28 pm PST
Subject: Re: code metaphor

[Martin Taylor 921120 18:00] (Bruce Nevin 921118 10:35:06)

I hope that this will provide enough fun to compensate for the use of net bandwidth on what is not PCT. A view over the fence.

> Our discussions over the past
>almost two years have not made it clear to me where you stand on
>the literal merit of the metaphor.
>
>To see that this metaphor is wrong you have to recognize that a
>code has no information structure in it other than that imparted
>to it by the process of encoding; that is, on this view an
>utterance can exhibit only such information structure as was
>previously present in the message or belief held to be "in" the
>sender.
>
>There is no warrant for claiming that perceptions associated with
>an utterance (as the meanings of the utterance) are structured in
>the way that the utterance is, and there is much to suggest that
>they are not.

I partially answered this (921118 13:30), but I think you would be amused by some of the argument currently going on about the meaning of "meaning" on the Usenet group sci.cognitive. If you think I lean toward "encodingism" look at these (unattributed to save reputations) quotes that I read today...

1.

Consider a theory analyzing meaning-preserving transformations. Let the said transformations be regarded as argument-valued functions (in the general sense, rather than taken in extension) on arguments, i.e. self-contained units of discourse. (An

elementary argument may be construed as a statement, i.e. a sentence in a context.) In order to get a reasonable theory of such beasts, you will want them to be broken down to the elementary level, where each transformation will correspond in extension to a one-one function $t:A \rightarrow A$. Of the said functions, one would clearly be the identity I_A ; moreover, if an elementary meaning-preserving transformation is characterized by a function $t:A \rightarrow A$, then it is to be expected that there exists another such transformation corresponding to its inverse, $t^{-1}:A \rightarrow A$. Finally, and most controversially from a natural language standpoint, the philosopher will expect that, given two meaning-preserving functions $t,s:A \rightarrow A$, their composition, $ts:A \rightarrow A$ will likewise be meaning-preserving. Given these conditions, which obviously likewise be meaning-preserving. Given these conditions, which obviously can be motivated in a way not depending on any reference to generative processes, all arguments will be partitioned by the equivalence relation induced by the class of all meaning-preserving transformations, into equivalence classes of synonymous arguments. Once your theory starts referring to the said equivalence classes by including them in the ranges of its existential quantifiers, you will thereby commit yourself to an ontology of abstract propositions and concepts.

If you still cannot understand what I mean, you will have a problem in any analytic discipline. Go study logic.

(Name deleted)

P.S. Is anybody out there objecting to transitivity of synonymy?

P.P.S. As for the person who intoned that meaning-preservation was a folk concept: if you really mean it, kindly go to a far-off place, and die by the Death of a Thousand Cuts. I can't mean that, can I? But to resolve any residual doubts, ask a Chinese friend to explain.

2.

>>

>>If you don't feel that being bound by
>>logic is a *moral* obligation, would you still have any basis for a
>>reasonable expectation that your tooting your car's horn at a bunch of
>>Samoan Hell's Angels does not *mean* your informed consent for them to
>>peel off your skin and rip out your heart? Do you have a glimmer of
>>understanding that your human rights are not worth the market value of
>>the chemicals that comprise your anatomy, *if* there be no such things
>>as matters of fact about meaning?

>

>It seems to me that this statement mixes several immiscible issues.
>Let us remember that human rights are based on a *commitment* to
>logic. We shouldn't believe that logic is fundamentally 'right' or something.

If you do not believe that it is fundamentally right, you are not committed to it.

3.

>Moreover, if understanding talk about "meaning" requires the adoption of
>some theory of meaning, then it would follow that philosophers can seldom
>understand each other when they talk about this subject, since there have
>been many and various theories of meaning. Thus, philosophers who hold
>different theories must be talking about different things when they discuss
>"meaning"; they only think they are talking to each other about the same
>thing.

I may be getting the wrong impression, but I do believe that philosophers more often than not do not understand each other when they talk about meaning. And then, after a while, you see them start talking only to those people who do understand them, and their theory.

>>I did NOT insist that there are reasons to formulate theories of meaning - I
>>took that for granted. ...

>I thought it might be interesting to ask you _why_ you take this for
>I thought it might be interesting to ask you _why_ you take this for
>granted. Is there some innate joy in formulating theories? Are they perhaps
>aesthetic objects that we admire for their beauty? Or does this theory have
>some _use_? If the answer is, "well, if you have to ask, then you can't
>understand", then I will, of course, have to desist from my questioning.
>Formulating theories of meaning is, perhaps, a natural urge that some
>people have and that can't be explained further. Those who don't get the
>urge can't expect to participate.

The answer is: the theory has to be there in order to explicate what we are really talking about. Some people may find some "innate joy" in formulating theories or contemplating their aesthetic value. I think however there is a more practical issue here.

4.

> P.S. Is anybody out there objecting to transitivity of synonymy?

I would object to the assumption that synonymy is a property of human language. Trivially, if we believe that two expressions "mean the same thing" then we have some meaningful criteria to distinguish them and some reason for doing so and, unless you want to multiply abstract entities endlessly, they therefore do not "mean the same thing". More concretely, for any pair of human language expressions that you claim to be synonymous in a given context, I claim that I can produce another context in which they are not synonymous. Finally, while genuine synonymes *might* exist as transient phenomena in human language, surely one would quickly be driven out by the other for simple reasons of efficiency (would you want to keep two identical copies of the same procedure in a program?).

Martin

Date: Fri Nov 20, 1992 9:32 pm PST
Subject: Science rejects are model -- unreviewed

[From Bill Powers (921120.2000)]

News flash. Science has returned the arm model article by Greg and me unreviewed. The explanation was that its priority was judged against all other materials currently on hand and found too low to be among the 20% they will publish. It took them three weeks instead of the customary two to reach this decision.

It will be interesting to see what is published in the last couple of weeks of February.

So where next?

Greg, I tried to call you but nobody's home.

Best, Bill P.

Date: Sat Nov 21, 1992 8:57 am PST
Subject: Using Standards

from Ed Ford (921121:0953) I said on (921116)

>>Any time I deal with anyone, whether in private practice or elsewhere,
>>standards and rules are a part of his life.

Rick Marken (921116.900)

>I just don't get it, Ed. What does "teaching standards" have to do
>with the PCT based view of human nature? What I get from PCT is the
>idea that nothing could be less important -- the actual substance of a
>persons references for relationships, program categories, principles,
>"standards", etc matters only in terms of how these satisfy higher
>level goals. The system should just be error free -- and this happens
>by having working (conflict free) control systems....

I am not teaching standards, but the intelligent evaluation and use of the ones people create for themselves. Or, I am trying to help people deal with the standards in the environment in which they find themselves to satisfy their own goals. An example would be helping a person to think through the best way to satisfy the goal of getting released from a lock up facility within the reality of his/her present environment. When trying to help others achieve or satisfy "these higher level goals" that you suggest, how do you think a person does this? When you're down in the trenches, you have to be very practical.

They can only achieve their goals by establishing in their own mind criteria (standards, rules, guidelines) upon which they are going to base their decisions. You say the system should be "error free" and I agree, and then you go on to say "this happens by having working (conflict free) control systems" to which I agree. My question is how do you help another system get to that point? When you set a reference for driving on the freeway, for establishing a closer relationship with a member of your family, for satisfying an employer or improving your job performance in a working environment, for employing a worker, or for just buying food at a grocery store, you sure do have standards or criteria based on your reference for the choices you're going to make to achieve your goal. And in order for you to function in the environment in which you find yourself in some of the above situations, you are going to have to become aware of what the agreed to standards or rules are that others in that environment have agreed upon to live by so that you and they can function cooperatively. You can't see a reference for safe driving but you can see stop signs, speed limits, and you are made aware when you get a state drivers license of the various rules or standards for driving.

Obviously, my friend, you have never read Freedom From Stress which goes into great detail on explaining the relationship of standards, principles, or whatever you want to call it, to the other levels on the hierarchy of control and how all of that understanding helps people control much more effectively and efficiently for references or goals. You just don't deal exclusively with the highest goal. And, more importantly, the various people with whom I work evidence a need for help in learning how to use their

system more efficiently so that they can function more effectively and get what they want.

People are able to articulate, prioritize, and evaluate not only references, but set appropriate standards or rules or criteria that will help them reach their goals. Also, these rules or standards will then act as a guide for the various choices they have to make IF they have learned to use their system properly. What I am trying to say is you teach people how to use their own system, to set their own goals, their own standards upon which they can make choices, because PCT teaches me that that is how the system is designed. This hierarchical system is highly inter-connected, cross-connected, interdependent, and being able to satisfy goals often demands the awareness and evaluation of all these various levels. And you know what, it all works.

This past week, I was at a conference where I was exposing educators from around the country to my Teaching Responsible Thinking card. What excited them most was the obvious connection between PCT and the practical use of the concepts on the card, regardless of the setting, whether home, school, business, or wherever. It gave them a very logical and clear understanding between PCT and what I was teaching them. And, they saw the practical, down-to-earth aspects of what I was saying.

You were absolutely wrong when you said...

>but you should be clear that this focuses your treatment strategy on
>getting a person to act in ways that are better FOR YOU -- and,
>incidentally, for the person him/herself.

I never, ever push people to act in a way that would be better for me. That is absolutely wrong. Please explain to me how this focuses my treatment strategy on my goals (except that of helping them to function more effectively and responsibly on their own). Have you ever seen me work with anyone or explain what I do through a role play demonstration. I suggest you read the role plays in chapters nine and ten in Freedom From Stress. If these people are a part of my life, a necessary part of the environment in which I attempt to live and work cooperatively with others (for example, at work or at home), I have to find out what their goals are, what they are planning to do, how they perceive things, so that I can deal with my life within the reality of the choices these others are making. In my counseling, it is the clients who are asking for help in learning how to deal with their world in such a way that they can satisfy their internal reference signals, including getting along with the people in their lives. It is these living control systems that are asking for help. They are asking to be taught the skills of functioning more efficiently and to learn how to reach their goals with the least amount of hassle. My goal is to help them with what they want. The last thing I do is to impose my values or beliefs on them. They are going to have to deal with the consequences that are a result of the goals, standards, and choices they make. I teach them how to manipulate themselves, to ride their own bike, to make their own choices, to satisfy their own goals. To manipulate people in such a way as to get them to do what I want is totally against good, sound counseling and teaching and totally against the PCT design. It is totally repugnant to everything I believe.

Ed Ford ATEDF@ASUVM.INRE.ASU.EDU

Date: Sat Nov 21, 1992 4:10 pm PST
Subject: pain & reorganization

[Avery Andrews 921122.1102] (Bill Powers (921120.0830))

>You're reinventing S-R theory. Your proposal is essentially how neural
>network people are trying to use neural nets -- to connect inputs to

I don't think so (and neither, it seems does Martin Taylor, who thinks I'm reinventing reorganization of perceptual input functions).

>The catch in this is "in such a way as to typically induce behavior that
>will alter the sense-vector" [in the right direction and by the right

Well, there is a bit of a catch there, but I think you've already come up with two ideas that would work some of the time. For an e-coli type creature, having the fear wire make a positive contribution to the tumble rate. For more complicated critters, the fear wire might make a positive contribution to the reorganization rate. There are doubtless other possibilities, but either of these ought to tend to keep the critter out of circumstances like those in which the pain wire fired.

Perhaps the idea could be described as one of ceating `derived intrinsic (negative, in this case) reference levels. Or have I missed something, and am sitting somewhere way out in left field?

(Martin Taylor 921120 11:30)

>This looks to me so much like Bill's reorganization of the perceptual
>input function that I can't tell the difference. I had originally thought

It would be, I think, that what I'm suggesting is a sort of one-shot `freezing' operation, which could be used to drive reorganization, & perhaps in other ways.

Avery.Andrews@anu.edu.au

9211D Done

Date: Sun Nov 22, 1992 9:36 am PST

[From Hank Folson (921122)]

From the Los Angeles Times, November 18, 1992: A researcher, Irene Pepperberg of the U. of Arizona has been teaching an African Grey parrot named Alex how to talk (i.e. communicate, not just mimic) for 15 years. The bird has a 90 word vocabulary and can label items by color, shape & material. When Alex compares 2 items, he can tell which is larger, and what attribute is the same or different.

Pepperberg teaches Alex the way she thinks the birds learn in the wild: While Alex watches, she teaches another person, rewarding or scolding them as they learn the new word. The process is then repeated directly with Alex. As a reward to reinforce (their word, not mine) the meaning, Alex is not given food, but is given the object he identifies to play with.

There may be a lot of language basics available from this research. Too bad they aren't using PCT as there are obvious connections and both the researcher & the bird could learn more.

Students of Language:

In a recent post on language, I noticed something missing from the post and the replies: No one said or asked what the speaker and listener were controlling for. What is even stranger to me is that the language posts in general, while clearly concerned about applying PCT principles to language, often approach the example sentences from the point of view of stimulus- response observers who are observing and completely accepting the behavior of the speaker and listener. Is there ever any discussion of the controlled variable(s)? Doesn't most of the discussion of these linguistic examples evaporate when you consider what the speaker is controlling for? The statements then usually make perfect sense as controlling actions, even though the language structure may leave a lot to be desired.

Mary Powers(921116):

Thank you for the excellent capsule summary of PCT research. It was so good I must remind the 'net that this was a test of a list of questions based on PCT principles to make it easier to initially evaluate theories. Any of you interested in PCT applications, please type up if you have a refinement or replacement for the list.

Best, Hank Folson

Date: Sun Nov 22, 1992 11:15 am PST
Subject: Re: Complex perceptions

[Martin Taylor 921122 13:40] (Bill Powers 921120.1300)

Bill,

One of us is missing, badly, a point the other is trying to make. Everything in your posting seems crystal clear to me, up to the point where you say:

>I think that the idea of complex perceptions, vectors that have
>meaning in themselves, is a vestige of an older way of modeling
>behavior based on a realist epistemology. We certainly experience
>multiplicity of perceptions. But until we form the right perceptual
>functions at a higher level, we do not perceive pattern in them, and
>none of the implicit patterns, infinite in number, has any effect.

If you are talking about "meaning in themselves" from the viewpoint of the person in whom the vector of perceptual signals exists, I quite agree. But remember that, right at the beginning of this discussion, we stipulated that the meaning of the vector is in the mind of the analyst, not of the perceiver.

I'm afraid I can't continue this here. I am spending more time editing out bad characters on the line than on the text I want to send. I'll send the rest when I get back to work.

Sorry. Martin

Date: Sun Nov 22, 1992 12:46 pm PST
Subject: Gloomy Sunday

[From Bill Powers (921122.1300)]

I'm disappointed that Science didn't want the arm paper. It was a good paper. They didn't realize what they were looking at, but that's nothing new. Greg and I will work on a longer version for submission elsewhere when the adrenaline comes back.

I think I'm going to concentrate on modeling for a while. The Primer series seems to be coming along; another one soon. Right at the moment I can't work up a lot of enthusiasm for hypothetical linguistics or abstract neural networks. It's back to spadework at the foundations for me. Let others worry about how to decorate the penthouse, should it ever be finished. Modelers of the world, arise!

Best to all, Bill P.

Date: Sun Nov 22, 1992 7:55 pm PST
Subject: privileged points of view; language experiments

[From Bill Powers (921122.2000)] Martin Taylor (921122.1340) --

>If you are talking about "meaning in themselves" from the viewpoint of the
>person in whom the vector of perceptual signals exists, I quite agree. But
>remember that, right at the beginning of this discussion, we stipulated that
>the meaning of the vector is in the mind of the analyst, not of the perceiver.

The analyst is a human perceiver. In my own mind, I do not allow privileged points of view in HPCT: whatever any human observer can do must, at least in principle, be explained by the same model. So if the analyst sees meaning or pattern in the vector, the analyst's perceptions are combining the individual elements of the vector and reporting a unitary perception, a scalar perception, in the manner I proposed. If the analyst's mind also makes use of information in the vector in the same way, the same model applies. There is no place to stand from which to see what the pattern "really" is. The same model must apply to the observed and observing human systems. It must also apply to a human being looking at a model, too.

The computations that produce the individual elements of the vector might well interact, so that the elements are not independent. That subject has to do with the internal design of one level in the hierarchy; what the next level makes of the vector is a different subject.

Hank Folson (921122) --

Your post is to the point. We have to consider what people use language for, not just try to explain language by examining words.

It would be interesting to devise some experiments to see why people use language as they do. I'm thinking of some task that requires two (or more) people to cooperate, but that requires them to convey information to each other. One way to test ideas about the elements of language would be to start with the most primitive operations that you think all languages contain, such as naming. Of course other primitive operations could be proposed, but the point is to test them to see if they really are primitive by seeing what CAN'T be done within the restrictions of the primitive elements.

Suppose that the people are allowed to send symbols back and forth by picking them out of a list. The symbols are arbitrary, not familiar words, and can be defined only through

common experiences with the task and with each other. How would people establish meanings for the symbols? What sorts of experiences would they attach symbols to first? How would they combine this process of building a language with learning to see the environment in mutually-consistent ways? How would they teach the use of the symbols to a newcomer?

While using a computer for this would make taking and recording data possible or at least a lot easier, tests of the idea could be done without special equipment. The list of symbols, for example, could be the 52 playing cards in a deck, laid out on a table. A symbol would be communicated by pointing to it or wiggling it (no normal verbal communication allowed). The common task could be assembling something, say a jigsaw puzzle. or it could involve one person telling another what to do in stacking blocks or cutting up pieces of paper or making brownies or anything else you can think of.

Some interesting problems come to mind. Attaching a symbol to an object might seem easy, but with no rules, how do you get across that you are attaching the symbol to the object and not to the actions by which you're calling attention to the object? How would you get it across that you're indicating some attribute of the object other than just its shape? In other words, how do you establish that the Jack of Spades means a corner piece in a jigsaw puzzle, and not the act of reaching out and picking up or pointing to any piece? Watching people work out problems like this ought to reveal a lot about how people invent languages. I would love to see our PCT discussions of language start incorporating experiments to see if explanations are right, instead of just plausible.

I think the deaf community would have a lot to teach about what happens when a group of people is brought together without any initial means of communicating. I claim that someone would ALWAYS invent a language, using whatever perceptions are easily manipulable as the tokens. In other words, I claim that it isn't language that's inherent, but ingenuity.

Best, Bill P.

Date: Sun Nov 22, 1992 11:08 pm PST

Subject: Re: privileged points of view; language experiments

(penni sibun 921122.2300)

[From Bill Powers (921122.2000)]

Hank Folson (921122) --

Your post is to the point. We have to consider what people use language for, not just try to explain language by examining words.

lots of us, both ``interactionists'' and more traditional psychologists, have been doing this for a long time.

It would be interesting to devise some experiments to see why people use language as they do. I'm thinking of some task that requires two (or more) people to cooperate, but that requires them to convey information to each other. One way to test ideas about the elements of language would be to start with the most primitive operations that you think all languages contain, such as naming. Of course other primitive

herb clark and gradstuds at stanford have a whole suite of such experiments. i'm not sure what they would consider primitive operations, if any, but i don't think naming w/b one. *referring* probably is (though of course referring needn't be verbal). for a fascinating read on this topic, i recommend clark and wilkes-gibbs, ``the collaborative construction of referring expressions'' (full ref. on request). referring is a dynamic and often volatile thing, whereas naming implies a mapping between stuff and symbols and has static connotations. i wouldn't be surprised if an adequate model of referring obviated any distinct model of naming.

Suppose that the people are allowed to send symbols back and forth by picking them out of a list. The symbols are arbitrary, not familiar words, and can be defined only through common experiences with the task and with each other.

why should we suppose this sort of artificial task will be revealing of anything particular to language?

Some interesting problems come to mind. Attaching a symbol to an object might seem easy, but with no rules, how do you get across that you are attaching the symbol to the object and not to the actions by which you're calling attention to the object? How would you get it across

there's literally tons of research on symbol manipulation. it's even been studied extensively in other species, including the usual suspects--chimps (see david premack, intelligence in ape and man) and dolphins (no handy refs., sorry). for that matter, one could cogently argue that skinner's pigeons can attach symbols to objects; we don't want to say this means that pigeons have human-like language, do we?

I think the deaf community would have a lot to teach about what happens when a group of people is brought together without any initial means of communicating.

i find this notion incoherent. w/ possible exceptions of severely damaged individuals (eg, those w/ autism, those severely abused in childhood), any group of individuals you bring together are going to have myriads of ways of communicating.

communication is a terribly important part of what human beings (and all organisms!) do. language may or may not involve symbol manipulation. but language is fundamentally communication. wouldn't it be a good idea to start studying a language at its foundations?

cheers. --penni

Date: Mon Nov 23, 1992 8:02 am PST
Subject: RE: Science rejects are model -- unreviewed

Bill and Greg [Re: WTP, 921120]

Incredible. Simply incredible. You must have had the misfortune to submit just at the time Bizzi sent them another manuscript. Is it the case that they publish his without review -- surely reviewers would not let them pass.

Are you considering Jean Arcady's journal, *Adaptive Behavior*? Or have you even had time to contact another? (Time for *Journal of Living Control Systems*?)

Tom Bourbon

Date: Mon Nov 23, 1992 8:54 am PST
Subject: Re: Knowledge Transfer

Cynthia:

Thanks for your interest in my work.

The model I am developing is largely macro. I draw heavily on a preliminary model of sociological self-organization proposed by Lee Freese (Washington State Univ.) in his 2-part article "Evolution and Sociogenesis" in the 1988 Advances in Group Processes pub. by JAI Press. Freese proposes a common unit for the analysis of change in ecosystems and human social systems: the resource transfer. Resources, as defined by Freese, may be tangible or intangible. Transfers are essentially one-way rather than reciprocal. His model directly links local action and higher-order "structure."

I try to elaborate an aspect of the model that Freese does not give adequate attention to: how transfer of knowledge claims from one interactant to another is a necessary to sustain larger (macro) resource configurations.

If this sounds interesting, you might want to read Freese. I find it difficult as yet to really discuss my work outside that contextual reference. However, I don't think that Freese or my own current work would be helpful to your current work, as neither of us goes in-depth into aspects of interpersonal communication.

However, Cecilia Ridgeway recently (with past year) published an article in the journal Social Forces that I think would be highly relevant to your work. I don't have the exact reference, because I have an advance copy of the draft. Let me know if you have any problems locating it and I'll get an exact reference for you.

Good luck on your writing!

Greetings, Mary Lee

Date: Mon Nov 23, 1992 8:56 am PST
Subject: Re: Science rejects are model -- unreviewed

[from Gary Cziko 9211212.0120 GMT]

Bill Powers (921120.2000) reported:

>News flash. Science has returned the arm model article by Greg and me
>unreviewed. The explanation was that its priority was judged against all
>other materials currently on hand and found too low to be among the 20%
>they will publish.

But didn't Science publish the Bizzi paper (or was that BBS)?

I can't help wondering how you and Greg identified yourselves on the submitted paper. Was it Powers, Independent Researcher, Durango, CO and Williams, Independent Researcher, Gravel Switch, KY? I am sad to admit that if I didn't know you guys and saw such authors,

I would not initially expect much of the paper. And if I had to decide to cut 50% or so out of the review process based on initial encounters, bye bye Powers and Williams.

Or do you at least give the CSG as your affiliation? Perhaps you need to have Bruce Nevin establish a P.O. Box for you in Boston or Cambridge to get a respectable address.

While I can understand yours and Greg's frustration, it really is an eye-opener in terms of how science in general (and Science in particular) operates. I hope you have better luck in Adaptive Behavior (or whatever the Jean-Arcady Meyer edited journal is called) or wherever else you decide to submit it.

--Gary

Date: Mon Nov 23, 1992 10:12 am PST
Subject: foundations of language

[From: Bruce Nevin (Mon 921123 09:58:11)]

No time to keep up, but as I sneaked a glance at the most recent this AM this caught my eye:

(Bill Powers (921122.2000)) --

```

/*****
*
* I think the deaf community would have a lot to teach about what happens
* when a group of people is brought together without any initial means of
* communicating. I claim that someone would ALWAYS invent a language,
* using whatever perceptions are easily manipulable as the tokens. In
* other words, I claim that it isn't language that's inherent, but
* ingenuity.
*****/
```

(I'm going to put quoted text in C-style comment boxes--my editor has a filter to do this--assuming that reformatting that would produce a bigger error for you than reformatting line-initial ">" to the interior of the paragraph.)

When has a deaf community invented a language in the absence of an environing community of people possessed of language who intercommunicate with the deaf folks by whatever means they are able?

You made a similar point, Penni. However, you went on to say:

Penni Sibun <sibun@PARC.XEROX.COM>

```

/*****
*
* communication is a terribly important part of what human beings (and
* all organisms!) do. language may or may not involve symbol
* manipulation. but language is fundamentally communication. wouldn't
* it be a good idea to start studying a language at its foundations?
*****/
```

I disagree that language "is fundamentally communication". Even more basically, people use language to create information that could not exist without language and could not have existed prior to language.

This is not to say that the perceptions ("meanings") with which words are correlated could not have existed, though some categories may well not be universal, but rather that higher-level organizations of those perceptions (which are themselves perceptions, of course) are not merely communicated by means of language, but are consensually created by means of language.

Nor is this to say that language is not used for communicating. But language is not really very good for communication, as witness your (Penni) twitting me about the nature of flirting. Of course talking about flirting is a far cry from doing it. And explaining a joke doesn't really "get" the humor of it, and so on. Just my point. Communication is carried out at a more basic level than language. We can use language for communication, but such use is no less semaphoric than the myriad other means we use for communication, what a modern Alexandrian Greek poet calls "the language of the turning eye". The agreements we can reach with language are more articulate than those typically reached by nonverbal means, but that is a function of the information-constituting capacity of language. The actual seeking for and communicating of agreement proceeds at a more primitive level irrespective of the specific content of the eventual agreement.

So I am arguing that communication is not at the foundations of language but rather antecedent to them. To be sure, communication has an important role in the story Harris proposes about the evolutionary origins of language (in _Language and Information_ and in _A Theory of Language and Information: a mathematical approach_). But to get at the foundations of language you must study something that is unique to language.

Just to head off a red herring, note that Shannon/Weaver communication theory is not about communicating in this sense either, any more than its near synonym information theory is about meaning in any general sense of the word.

Gotta run.

Bruce bn@bbn.com

Date: Mon Nov 23, 1992 12:31 pm PST
Subject: Re: Re: Gloomy Sunday

From Tom Bourbon [922311 -- 11:45 CST]

My apologies for the aborted message [922311 -- 8:20 CST]. I was hoping mad after reading Bill's posts concerning the rejection from *Science* (as in BIG SCIENCE). I realized my reply was going to turn into a flame job and tried to abort -- that would be CNTRL-C. Instead, I hit CNTRL-Z and immediately logged out. Imagine my surprise when I found the half-line message from me, sitting in the mail box.

"Hoping mad." What a slip! No hope, but lots of hopping!
Tom

Date: Mon Nov 23, 1992 1:09 pm PST
Subject: Ms communication; Science

[From Bill Powers (921123.1000)] Penni Sibun (921123.2300) --

>> Your post is to the point. We have to consider what people use
>>language for, not just try to explain language by examining words.

>lots of us, both ``interactionists'' and more traditional
>psychologists, have been doing this for a long time.

Funny, Mary predicted almost exactly what you would say.

It's strange how everything I suggest, working out of the PCT framework, turns out to have been done or suggested by someone else in some branch of linguistics or psychology. Of course these things have been done by people who violently disagree with each other on almost every fundamental issue about language and psychology, and there seems to be no foundation of accepted principles in either discipline.

I detect a pulling together of the wagons in the face of outside threats.

>>It would be interesting to devise some experiments to see why
>>people use language as they do. I'm thinking of some task that
>>requires two (or more) people to cooperate, but that requires them
>>to convey information to each other.

>herb clark and gradstuds at stanford have a whole suite of such
>experiments. i'm not sure what they would consider primitive
>operations, if any, but i don't think naming w/b one.

How about telling us what experiments they did and what they have discovered? This would seem to be an example of what I was asking about a few days ago, a base of non-controversial observations on which all linguists could agree, right? When you say they have a whole suite of "such experiments," do you really mean experiments of the kind I have in mind? Or are you just free-associating?

>>Suppose that the people are allowed to send symbols back and forth
>>by picking them out of a list. The symbols are arbitrary, not
>>familiar words, and can be defined only through common experiences
>>with the task and with each other.

>why should we suppose this sort of artificial task will be
>revealing of anything particular to language?

But I thought you just said this had already been done, with whole suites of experiments of the kind I am proposing. Perhaps you don't understand what I am proposing. I didn't mean that the passing of symbols would be the task; that would just constitute the means of communication. The task could be anything at all, as I said later in the post, such as teaching someone to make brownies.

>> Some interesting problems come to mind. Attaching a symbol to an
>>object might seem easy, but with no rules, how do you get across
>>that you are attaching the symbol to the object and not to the
>>actions by which you're calling attention to the object? How would
>>you get it across

>there's literally tons of research on symbol manipulation.

Good. So what facts have been found concerning the question I raised above, about how you get across the nature of the perception to which you're trying to attach a symbol? Or has the research all been too crappy to answer a specific simple question like that? "Tons of

research" doesn't mean tons of _good_ research. If you do good research, you don't need tons of it.

>one could cogently argue that skinner's pigeons can attach symbols
>to objects; we don't want to say this means that pigeons have
>human-like language, do we?

Is Skinner's research now accepted in linguistics? Anyway, I'm not much interested in what one could "cogently argue." Cogently arguing is a dilettante's game. I'm more interested in conclusions that explain good data and that smart people can't think of any alternatives to.

>>I think the deaf community would have a lot to teach about what
>>happens when a group of people is brought together without any
>>initial means of communicating.

>i find this notion incoherent. w/ possible exceptions of severely
>damaged individuals (eg, those w/ autism, those severely abused in
>childhood), any group of individuals you bring together are going
>to have myriads of ways of communicating.

So there was no motivation for deaf people to learn ASL, because they already had ways of communicating? You find this notion _incomprehensible_, Penni. There's a difference between your lack of comprehension and the incoherence of other people's notions. It seems to me that you're trying harder to show that you already know everything worth knowing than to remedy your lack of comprehension.

>communication is a terribly important part of what human beings
>(and all organisms!) do.

Then why is it so hard to communicate with you?

Gary Cziko (9211212.0120 GMT) --

>But didn't Science publish the Bizzi paper (or was that BBS)?

It was Science. In September of 1991 I wrote a technical comment on the Bizzi paper, which was rejected. All three of the reviewers, one of whom was Bizzi, rejected the commentary and said that I should first prove that my model does what I say it does.

This morning, at Mary's brilliant suggestion, I wrote letters to Bizzi and (via Science) the two anonymous reviewers. I explained to Science that I wanted these reviewers at least to see the model, as they had all expressed doubts about it. In the cover letter to Science I also quoted the relevant parts of the 1991 reviews, as follows:

Reviewer 1 said

"Although I agree with the general point of his [my] comment, I feel that it represents mainly a misunderstanding of ... Bizzi. I would very much like to see the development of a more general model, one that incorporates the basic observations of Bizzi et. al. and the philosophy of Powers ... but I would rather see the Bizzi et. al. point of view countered with data and research."

Reviewer 2 said

"Finally, the main point of Power's [sic] technical comment seems to be an advertisement for his model ("Soon to appear in a journal in your neighborhood.") Certainly, if his model can do all he claims, it would be of interest, but too much of this technical comment is based on acceptance of Power's stipulation that his model, offered for mid-1992 delivery, works. However, in the absence of detail, his comments fall flat."

Emilio Bizzi (the target of the technical comment) said

"In summary, we believe the burden of proof is on Dr. Powers to show that it is indeed simple for the CNS to generate posture and movement using a feedback control scheme which resides predominantly in the CNS and which must combine the diverse sensor information sources he describes. Such a scheme must surely have to take into account the various different feedback loop delays, and the complexity of muscle, sensor and limb dynamics."

I corrected the mistaken reference to "Dr."

I did not say or imply that Science should reconsider the Arm article. We will see what happens.

Best to all, Bill P.

Date: Mon Nov 23, 1992 2:34 pm PST
From: mmt MBX: mmt@ben.dciem.dnd.ca
TO: * Dag Forssell / MCI ID: 474-2580
Subject: CSGintro

Dag,

I was looking over your intro dated Nov 18, which I printed out, and I noted that all the references to text files and programs seem to have a form that suggests IBM formats (file.exe and the like). I realize that Bill's demo programs are restricted to IBM compatibles, but the text files should not be. Is my interpretation of the document correct, and there is no way of getting introductory documents if you don't have an IBM-compatible machine?

If so, I think it should not be so, and if not, the document should tell those of us with good machines how we can use them for the good of our souls.

Martin

Date: Mon Nov 23, 1992 3:23 pm PST
Subject: Re: privileged points of view; language experiments

[Martin Taylor 921123 11:40] (Bill Powers 921122.2000)

Continuing my interrupted posting of 921122.1340, to which Bill has commented:

>The analyst is a human perceiver. In my own mind, I do not allow
>privileged points of view in HPCT: whatever any human observer can do
>must, at least in principle, be explained by the same model. So if the
>analyst sees meaning or pattern in the vector, the analyst's perceptions
>are combining the individual elements of the vector and reporting a
>unitary perception, a scalar perception, in the manner I proposed.

>...
> There is no place to stand from which
>to see what the pattern "really" is. The same model must apply to the
>observed and observing human systems. It must also apply to a human
>being looking at a model, too.
>
>The computations that produce the individual elements of the vector
>might well interact, so that the elements are not independent. That
>subject has to do with the internal design of one level in the
>hierarchy; what the next level makes of the vector is a different
>subject.

To all of which, I naturally say, "of course." But I suppose there might be listeners to whom it is not obvious, and therefore it is a worthwhile expansion of something I thought not worth mentioning.

Now let's continue with where I think one of us is missing the other's point, and, not having a privileged point of view, I can't tell whether it is you or me. So I perhaps have to spell out more than I think should be necessary.

(Bill Powers 921120.1300)

>You replying to Rick:

>

>>>The question in my mind is: why should it be true that
>>>at some point, some resource (a cell or cell assembly) is going
>>>to have to represent the state of the distributed representation
>>>to someone other than the person watching the model -- ie. there
>>>will have to be a perceptual signal that represents the state
>>>of the distributed representation...

>

>The "someone else" is a higher level system. This discussion is
>making me go back through some old musty trains of thought that
>eventually led to HPCT.

The way I put it may have been misleading.

On reflecting over posts by you and Rick, I think there are three separable reasons for considering these complex percepts, even while recognizing that, in themselves, they have no reality for the perceiver:

(1) as a label that we as analysts may use as a shorthand to communicate amongst ourselves about groups of things that may be seen to be happening together (in us, the patterns are the perceptual input functions for scalar perceptual signals, but we are talking about complex things happening outside of us).

(2) to represent a subset of the different perceptual signals being controlled at one level, many or most of which may be inputs to several different perceptual functions at the next higher level, and that therefore tend to work together as a modular group (even if the modules do tend to overlap).

(3) as a base on which higher-level ECSs can develop by whatever means of reorganization creates and effectively links new ECSs (from a bottom-up point of view, this is a kind of non-linear principal components analysis, but executed, we believe, by random-walk processes rather than algorithmically).

>The computations that produce the individual elements of the vector
>might well interact, so that the elements are not independent. That
>subject has to do with the internal design of one level in the
>hierarchy; what the next level makes of the vector is a different
>subject.

The next level probably makes many things of the vector, but an outside analyst would say that those things were of a kind different from the elements of the vector, or from the vector itself. Internally, of course, they aren't--we use the simplifying concept of "neural current" or "neural voltage" to signify what carries the signals, and all of them are of the same kind. What is of different kind is the CEV represented in the controlled signal. In a structured perceptual signal, there are many CEVs, and thus many degrees of freedom (or perhaps I should put it the other way round).

Let us assume that all the simple perceptual signals in the complex are under good control--the errors are all close to zero. But the world provides disturbances. and those errors (momentarily) deviate from zero. If (as an outside godlike observer might see) two or more of these transient errors tend more often than not to be in the same direction and in the same ratio, it might often be more efficient for the organism to have an ECS that directly controlled the sum, perhaps as a substitute for the individual ones, perhaps as a higher-level ECS that used the individual ones as input. This is one version of "leaning on the world." The correlations among the perceptual signals in a structured percept would be due not to the structure of the percept, but to the disturbances generated by the world. If one ECS can control against a substantial part of the disturbances that actually occur, it is presumably efficient that such an ECS would be developed through reorganization.

We have argued (actually Bill pointed out) that reorganization tends to drive ECSs into controlling nearly orthogonal perceptual input functions, but we also know that there are far too many sensory degrees of freedom for this to happen overall. In the context of this posting, what this amounts to is a suggestion that any higher-level ECSs that take the signals of a structured percept as inputs would tend to develop into a principal components representation of the structured percept; they would tend to maximize the information transfer while minimizing resource usage.

When I talk about a "belief" as a structured percept, I am taking account of the multidimensionality of the perceptual structure. Many parallel perceptual signals may be components of a "belief," even if there is no actual control of the belief as such--no single perceptual input function to which it corresponds. There is only control of the multiplicity of scalar perceptual signals, all of whose related error signals are to be brought near zero. It is not necessarily satisfactory for any single function of those signals to be brought to a reference value, though it could be satisfactory for several orthogonal functions of those signals simultaneously to be brought to their several reference values. But then we would be talking about a higher-level Structured Control System.

Bill, I realize you are getting bored with "abstract neural networks." Is this the kind of thing you mean? I hope not.

Martin

Date: Mon Nov 23, 1992 3:41 pm PST
Subject: "Science" and Bizzi

[from Gary Cziko 921123.2100 GMT] Bill Powers (921123.1000) informed us:

>This morning, at Mary's brilliant suggestion, I wrote letters to Bizzi
>and (via Science) the two anonymous reviewers. I explained to Science
>that I wanted these reviewers at least to see the model, as they had
>allexpressed doubts about it. In the cover letter to Science I also quoted
>the relevant parts of the 1991 reviews, as follows: . . .

So it turns out that writing the technical review first of Bizzi means that your foot is still (slightly) in the door. There may be a lesson here for anyone wishing to publish in Science

You have put Bizzi and your two reviewers in very interesting situations. Will Science even forward your paper to the two reviewers? How will Bizzi react now that you and Greg have put your model where your mouths are? My prediction will be that the two reviewers (if they ever get to see the paper) will be positively impressed. I also predict that Bizzi will do his best to find all sorts of reasons why your model may APPEAR to work but is invalid due to what he will claim are certain simplifying assumptions that you and Greg made. What else can he do?

Do keep us informed as to how this plays out. This is the interface of normal science and revolution. And we on CSGnet are very fortunate to have first-row seats as the drama unfold (except for those of us who are also on the stage).--Gary

Date: Mon Nov 23, 1992 3:49 pm PST
Subject: Re: Complex perceptions

[Martin Taylor 921123 16:00] (Bill Powers 921120.1300)

This should have been part of my earlier posting of today...

>The HPCT view is that the pattern is imposed from above. When a
>higher perceptual function becomes organized, it receives signals
>that are lower-level perceptions, puts them through some sort of
>computing function, and delivers a perceptual signal carrying the
>scalar value of that function.
>...

>We do not learn to perceive patterns that are "really there." We
>learn to perceive patterns that are useful to control. Somewhere,
>way off in the philosophical distance, there may be a deep principle
>involved, to the effect that the patterns that prove to be the most
>useful to control, in terms of the effect of doing so on critical
>variables, happen to relate in some way to something that is in some
>sense "really there." But organisms do not know what is really
>there.

The second of these quotes leaves the first as something of a mystery. How can an ECS be developed that controls something "useful to control," if, at the same time, "the pattern is imposed from above." I have assumed that the perceptual signals being controlled have developed through some form of random reorganization, whereas "imposed from above" seems to suggest that there is some kind of direction or supervision involved. But the hint of that suggestion caused you to take off against Avery the other day. So there is a bit of bewilderment here.

My way of looking at it is to say that the second of these quotes is correct. We learn to perceive only that which we learn to control, and though we may at some point learn to control for many things, we retain those which it is useful to control. What we control against is as important as what we control for, in the sense that reorganization of the output connections allows effective, resource-limited actions that (usually) counter the disturbances that affect the perceptions we are controlling for. Different kinds of reorganization probably apply to the perceptual input functions and to the output connections, but both, presumably, must operate without magical knowledge of their optimum form.

Correlations among the disturbances imposed by the world provide opportunities for increasingly effective ECSs. Take those "correlations" in a more general sense, and you get the opportunity for different kinds of perceptual input functions; these "kinds" define the eleven levels of the hierarchy. As you said in your example of the balloon being blown up and then bursting, as seen by an organism with no transition level: "This new signal is the sense of transition, of growing. It has no other character."

But never would I say that these patterns are "imposed from above." The world provides the opportunity for higher-level patterns to be developed out of structured perceptions, and this opportunity comes from below. Only when the pattern has been developed into an ECS in control of its higher-level perceptual signal is there any sense in which the pattern is "imposed from above." It is in the operation of a mature hierarchy, not in its growth.

I have sometimes talked about the "mirror world" of perceptual signals and CEV's (Complex Environmental Variables). A perceptual input function cannot be stable unless the corresponding CEV has some (perhaps temporary) stability. It is not "useful to control" something that isn't there the next time you provide the pattern of outputs (actions) that have been reorganized. Even though actions are not controlled, they are another word for the outputs of an ECS, as seen by an outside observer. And if the actions sometimes lead to negative feedback, sometimes do nothing, and sometimes lead to positive feedback, that perceptual input function is not (at present) useful to control. Something needs further reorganization, and it may well be that the CEV just is not a stable aspect of the outer world. Then the perceptual input function will have to change before any control can be maintained. But this does not mean that the signals (perceptual signals of lower levels) that enter the perceptual input function also have to be changed. They, by assumption, are being very well controlled, thank you very much.

>I think that the idea of complex perceptions, vectors that have >meaning in themselves, is a vestige of an older way of modeling >behavior based on a realist epistemology.

Perhaps you can see that I am not talking about vectors that have meaning in themselves within the perceiving organism, whether that organism be the analyst or the analysand.

Do I still miss some point you are making? Or do you get my point any differently from before? Does it matter?

Martin

Date: Mon Nov 23, 1992 4:23 pm PST
Subject: Distributed control and G-ma signals

[From: Oded Maler 921123]

First I owe an answer about the possibility of distributed control which looks "as-if" it is hierarchical. Suppose x and y are the elementary perceptual variables and that the higher level variable is $z=ax+by$ (assume $a=1$ and $b=-1$ in the following example). According to the difference z^*-z , the higher system sets references for x and y , which essentially are of the form "increase" or "decrease" according to the sign of a , b and z^*-z . Graphically, if everything works well (we are in the range of controllability of x and y , etc.) our system will lead us from every point in the x/y plane toward the diagonal line representing $ax+by=z^*$. An observer of this system will certainly say that z is controlled.

Now consider the case when there is no explicit signal representing z , but rather two low-level ECS controlling independently for x^* and y^* , such that $ax^*+by^*=z^*$. Graphically the system will move from every point to the point (x^*,y^*) which is one point on the diagonal. An external observer who sees only z , will again conclude that z is a controlled variable, although it is not explicitly represented inside the system.

Which is better? It depends on the simplifying assumptions. If you assume that the controllability range of x and y are unlimited, and that z corresponds to some really-important CEV, then the hierarchical mode is better because it takes the shortest path to the diagonal, while the distributed control system will try to reach (x^*,y^*) from everywhere, even from other points (x',y') which satisfy $ax'+by'=z^*$. On the other hand the hierarchical control scheme does not "care" about pushing x and y into extreme values beyond their capabilities. If x and y reflect some important EVs, this might be a problem. (My association are soldiers dying for their country, or if you prefer, Lymphocytes dying for their body).

Last year Bill demonstrated that under certain conditions hierarchical control can be achieved of two higher-level perceptions e.g., $z=x+y$, $w=x-y$. Again, I think that under the same assumption you can achieve a distributed version.

I don't have any ideological bias against "signal for every percept", and hierarchy is very useful (when it works), especially when the reference signals for the higher variables may change, but I just wanted to demonstrate that a distributed control system can fool an external observer and look like a hierarchical one.

Consider now the "special status" of those perceptual variables which have a "dedicated" scalar signal. If you take their neural correlates to be currents on some neuronal wires, you may have some problem to explain the development of new "concepts" and any other "learning" phenomena in adults, because according to the current dogma in neurobiology (please correct me if I'm wrong) the changes in the connectivity patterns of neurons cease at some stage, so how are new controlled perceptual variables acquired? What is the perceptual "meaning" of a signal on some wire before it became dedicated to some variable? This is why I think one should not be so dogmatic about the actual realization of variables, and be ready to accept more sophisticated forms using which virtual signals can exist.

(I don't think that the question of all-or-none vs. degree, is relevant here, because both can be realized either by a scalar or in a distributed manner.)

Bill & Science: It takes an enormous amount of time to develop the perceptual capability to tell geniosity from charletanism and crackpothood, and I'm sure the editors of Science don't have the time nor the real motivation. If I were you, I would, in parallel to the development of a full version for Biol. Cyb., send the rejected stuff to Nature (it's already in a similar format, so it costs you nothing).

Best regards --Oded

Date: Mon Nov 23, 1992 4:40 pm PST
Subject: Re: Fowl Language, Fowl language

I'm in the new office. My new e-mail address does not seem to work (I didn't get subscribed to the list) but I've found that I can still access the net via my "old" address -- so I will continue using that.

Bill Powers and Greg Williams --

Sorry to hear about the Science article. Welcome to the club.

Hank Folson (921122) --

>Students of Language:

>

>In a recent post on language, I noticed something missing from the post and
>the replies: No one said or asked what the speaker and listener were
>controlling for.

Excellent point, Hank. I completely agree; it doesn't make much sense to try to apply control theory to a phenomenon if you don't know what variable(s) are being controlled; my mantra once again: Control theory is about CONTROL. If language does not involve control then control theory does not apply.

Richard S. Marken

Date: Mon Nov 23, 1992 4:44 pm PST
Subject: Usenet (NetNews) Mystery

[from Gary Cziko 921123.2223 GMT]

I occasionally check the Usenet (NetNews) version of CSGnet (where we are group "bit.listserv.csg-1") to see if it is keeping up with all the activity of CSGnet. I have no idea of how many people participate on CSGnet via Usenet, but it does offer some advantages over the LISTSERV format, particularly for those users with limited mail storage (other it does seem to run several hours behind the LISTSERV setup).

To my horror, I discovered that when I checked CSGnet through my local Usenet node there was not a trace of Bill Powers! Among messages dating back to 921117 there Marken, Taylor, Bourbon, Ford, Nevin, Jackson, etc. etc., and even Cziko. But nowhere was a post from Bill Powers to be seen.

All mail sent to CSGnet is supposed to be automatically sent to the Usenet new group. How mail from Bill Powers is being selectively discriminated against is beyond my comprehension.

I would greatly appreciate if others on CSGnet who have access to Usenet (NetNews) check out "bit.listserv.csg-1" and let me know via direct mail whether they can find Bill Powers or if anybody else appears missing. Hopefully, I will learn that there is something quirky about my local machine and Bill is getting out on other Usenet systems.

--Gary

Date: Mon Nov 23, 1992 5:12 pm PST
Subject: Server documents

[From Dag Forssell (921123.15.00)] Martin today - direct

>I was looking over your intro dated Nov 18, which I printed out, and I
>noted that all the references to text files and programs seem to have
>a form that suggests IBM formats (file.exe and the like). I realize that
>Bill's demo programs are restricted to IBM compatibles, but the text
>files should not be. Is my interpretation of the document correct, and
>there is no way of getting introductory documents if you don't have an
>IBM-compatible machine?
>
>If so, I think it should not be so, and if not, the document should tell
>those of us with good machines how we can use them for the good of our
>souls.
>
>Martin

I am taking the liberty of posting this so Bill Silvert and others can see it, since I can't quite answer.

I have tried to keep the document to the point: "A number of ASCII documents....are available...". I am suggesting to download the csg/Index. Have you done that? Several of the subdirectories are labeled: document. Does that answer your question? I have suggested the addition of comments to indicate ASCII, for this reason. Would that make a difference? (That is not up to me and I am no longer pushing for that).

I can't quite see (with my IBM brain) where your impression comes from. ASCII does not mean IBM, does it? Do other Mac enthusiasts have the same perception? Please make a specific suggestion for future editions if this is a real problem and not a fleeting impression.

Best, Dag

Date: Mon Nov 23, 1992 5:13 pm PST
Subject: foundations of language (B nevin 921123 09:58:11)

[From gabriel 921123 17:11 CST]

>[From: Bruce Nevin (Mon 921123 09:58:11)]

>No time to keep up, but as I sneaked a glance at the most recent
>this AM this caught my eye:

>(Bill Powers (921122.2000)) --
>/*****
>* *
>* I think the deaf community would have a lot to teach about what happens *
>* when a group of people is brought together without any initial means of *
>* communicating. I claim that someone would ALWAYS invent a language, *
>* using whatever perceptions are easily manipulable as the tokens. In *
>* other words, I claim that it isn't language that's inherent, but *
>* ingenuity. *
>*****/

>When has a deaf community invented a language in the absence of
>an environing community of people possessed of language who
>intercommunicate with the deaf folks by whatever means they are
>able?

Don't fish communicate non verbally? e.g. the UK stickleback, or lots of tropical species while nesting, to say nothing of Bluegill?

>You made a similar point, Penni. However, you went on to say:

>Penni Sibun <sibun@PARC.XEROX.COM>
>/*****
>* *
>* communication is a terribly important part of what human beings (and *
>* all organisms!) do. language may or may not involve symbol *
>* manipulation. but language is fundamentally communication. wouldn't *
>* it be a good idea to start studying a language at its foundations? *
>*****/

>I disagree that language "is fundamentally communication". Even
>more basically, people use language to create information that
>could not exist without language and could not have existed prior
>to language.

I'm really an amateur at this game, but it seems to me that a Shannon channel certainly communicates what the probable transitions of the source are. What this means depends on context. Many of the things that interest me seem to be descriptions of traversals of real or hypothetical decision trees, the real traversals often being to solicit the an estimate of probable truth of a proposition determining which branch to take at the sender's current node. The hypothetical decision trees (penni's very good example - X's room is down the hall on the right) being to establish context about an "aforesaid" which can henceforth be abbreviated by the symbol sequence "X's room" (avery's comment). Neither Chomsky or any of the programming language people like myself have any trouble with this. Question - probably trivial to answer - are the programming languages of interest to linguists as languages for communication with machines? I'm reminded of a comment of mine from years ago - a program is a contract between people and machines specifying what each must do to achieve some result desired by people.

>This is not to say that the perceptions ("meanings") with which
>words are correlated could not have existed, though some
>categories may well not be universal, but rather that
>higher-level organizations of those perceptions (which are
>themselves perceptions, of course) are not merely communicated by
>means of language, but are consensually created by means of

>language.

One of the nice things about programming languages is that there are ANSI standards for their meanings and working models for their understanding by machines. Yes, they are truly dull for linguists, but perhaps looking at a context free grammar might simplify the discussion because the meaning of a term cannot change subtly in the middle of a paragraph - or perhaps it can, but only in well defined ways. A dumb question perhaps, but is it worth thinking about? For example the elementary concepts of LISP 1.5 (the undefined terms of` axiomatic theories) can be used to build artifacts having apparently complex behaviours - e.g. LIZA. Are these as worthy of study as the PCT ARM, and if so why (or if not why). And any real light PCT could shed on the process of making good software will make its discoverer rich and famous.

>So I am arguing that communication is not at the foundations of
>language but rather antecedent to them. To be sure,
>communication has an important role in the story Harris proposes
>about the evolutionary origins of language (in Language and
>Information and in A Theory of Language and Information: a
>mathematical approach). But to get at the foundations of
>language you must study something that is unique to language.

>Just to head off a red herring, note that Shannon/Weaver
>communication theory is not about communicating in this sense
>either, any more than its near synonym information theory is
>about meaning in any general sense of the word.

Again, perhaps a dumb comment. The original Shannon Weaver paper was a very abstract affair. When I look at either ENIGMA (the biography of Turing, or the account of Bletchley Hall 1938-1945) or back issues of CRYPTOLOGIA, these make me inclined to argue that the red herring is not perhaps as much of a false trail as you might think. But it seems clear to me that there are subtleties in bruce's view of language I will not appreciate until we have more shared experience. Some of the recent articles about Rosetta Stone like problems that are still open might be of interest to linguists, and have at least to me a VERY Shannon flavour. Bruce will know of Bill Woods I am sure from Bill's long association with BBN. How much difference really is there between Shannon's traversals of a Markov Diagram, and an individual's traversals of one of Bill's semantic nets, or successive retrievals from penni's KB?. Or for that matter the successive retrievals from a Kanerva database as the algorithm seeks the closest match to a given pattern - Well, perhaps that's a bit far fetched. But how about two very similar Kanerva Databases, and the mapping between them by retrieval of closest matches to a succession of different patterns!! Perhaps that question IS worth some thought.

With Greetings to all, and malice to none. Avery's excerpt from another SIG suggests that our flames are low temperature. Perhaps the diversity of interests here is a damper for the free radicals of flame propagation. No pun intended.

Well, I probably need to find time to read Harris as bruce suggests, before succumbing to fatal hoof in mouth disease.

John Gabriel (gabriel@eid.anl.gov)

Date: Mon Nov 23, 1992 6:05 pm PST
Subject: Re: Ms communication

(penni sibun 921123.1600)

[From Bill Powers (921123.1000)]

Penni Sibun (921123.2300) --

>> Your post is to the point. We have to consider what people use
>>language for, not just try to explain language by examining words.

>lots of us, both ``interactionists'' and more traditional
>psychologists, have been doing this for a long time.

Funny, Mary predicted almost exactly what you would say.

it's a good thing we're all as predictable as we are, or we'd have even more trouble understanding each other!

It's strange how everything I suggest, working out of the PCT framework, turns out to have been done or suggested by someone else in some branch of linguistics or psychology.

it's not, of course. it's a major problem that *every* academic faces these days--there's far too much work out there already done for anyone to keep track of it. and, if any of us is going to have time to *think* at all, we have to deliberately lose track of some of it. on the other hand, if one has acquaintances in other disciplines, one has somewhat more breadth.

Of course these things have been done by people who violently disagree with each other on almost every fundamental issue about language and psychology, and there seems to be no foundation of accepted principles in either discipline.

i think ling and psych are far too broad to have single foundations of accepted principles. i think it w/b unhealthy if they did!

I detect a pulling together of the wagons in the face of outside threats.

well, don't infer that from me. i don't have even one wagon!

look. i think pct is interesting enough and fresh enough that someday it might make it into the canon and textbooks. it would break my heart to see it do so burdened by the same old tired old ideas about language (eg, that it's based on symbols) that have been sitting around for a long time doing not much work.

let me try to explicate distinctions that i previously was implying.

communication--something that most organisms (can we agree on at least all vertebrates?) do w/ each other a lot, in many different ways.

language--a kind of communication that is particularly flexible and complex, and seems to incorporate systematic ways for communicating about things not readily accessible in the here-and-now (including ``abstract'' things).

symbol manipulation--a practice of mapping bet. 2 sets of individuated things, in which members of the first set, and relations bet. them, are regarded as systematically

substitutable for members and relations in the second set. (sorry--this isn't a very good defn; i hope it's not impenetrable.)

i look on language as a specialized form of communication. i do not think that as it is used it involves symbol manipulation, at least in any straightforward way, such as is implied by ``naming.``

>herb clark and gradstuds at stanford have a whole suite of such
>experiments. i'm not sure what they would consider primitive
>operations, if any, but i don't think naming w/b one.

How about telling us what experiments they did and what they have discovered? This would seem to be an example of what I was asking about a few days ago, a base of non-controversial observations on which all linguists could agree, right?

well, clark & co are psychologists, and their observations are not non-controversial. but then again, i don't think anybody's are. at any rate, their results are of the statistical sort that seems to get rick upset.

When you say they have a whole suite of
"such experiments," do you really mean experiments of the kind I have in mind? Or are you just free-associating?

well, i certainly can't read your mind. you did however say:

>>It would be interesting to devise some experiments to see why
>>people use language as they do. I'm thinking of some task that
>>requires two (or more) people to cooperate, but that requires them
>>to convey information to each other.

and clark & co's work fit this bill.

>>Suppose that the people are allowed to send symbols back and forth
>>by picking them out of a list. The symbols are arbitrary, not
>>familiar words, and can be defined only through common experiences
>>with the task and with each other.

>why should we suppose this sort of artificial task will be
>revealing of anything particular to language?

But I thought you just said this had already been done, with whole suites of experiments of the kind I am proposing.

you've restricted yr specification to symbol manipulation, whereas yr proposal was to investigate language. from the papers i have on hand and what i remember, clark & co never insisted subjects communicate only via unfamiliar arbitrary symbols. their typical paradigm is to have two subjects who can hear but not see each other, and who have two, perhaps identical, sets of tools (or shapes or something) in front of them. one knows what they're supposed to build (usu. by having a model in front of her) and the other has to do the building. the frobs in front of them are not things which have names (or, often, easy description). the locus of study is how they use language to get the task done. at least half a dozen gradstuds have gotten or are getting theses out of this paradigm. it's a very rich mine.

>there's literally tons of research on symbol manipulation.

Good. So what facts have been found concerning the question I raised above, about how you get across the nature of the perception to which you're trying to attach a symbol? Or has the research all been too crappy to answer a specific simple question like that? "Tons of research" doesn't mean tons of good research. If you do good research, you don't need tons of it.

i agree with you. i imagine for a diff. reason, though: i think the research is misguided. sure, you can hold up an apple, and i can say, because i learned to, ``3 of clubs.'' but i don't think this is any more fundamental a skill than your saying ``apple'' and my saying ``a p p l e.''

unfortunately, my intro psych book has gone missing. and maybe my copy of premack is still back in ct. but i'll take an example from the latter. the chimp sarah was taught to associate colored plastic tokens with various things. she was quite good, for example, at knowing that red-square was associated w/ apples. she also knew that black-triangle--blue-circle--red-square was associated w/ her eating an apple in a way that no other configuration of tokens was. what was touted as particular proof of her using the tokens in an arbitrary fashion was that she knew that green-ellipse was associated w/ red things (and not green ones; chimps have color vision).

>one could cogently argue that skinner's pigeons can attach symbols
>to objects; we don't want to say this means that pigeons have
>human-like language, do we?

Is Skinner's research now accepted in linguistics?

i don't know. i was doing a reductio ad absurdum.

>>I think the deaf community would have a lot to teach about what
>>happens when a group of people is brought together without any
>>initial means of communicating.

>i find this notion incoherent. w/ possible exceptions of severely
>damaged individuals (eg, those w/ autism, those severely abused in
>childhood), any group of individuals you bring together are going
>to have myriads of ways of communicating.

So there was no motivation for deaf people to learn ASL, because they already had ways of communicating? You find this notion incomprehensible, Penni.

perhaps my referent was sloppy: by ``this notion'' i meant ``[bring] together [a group (of people)] w/o any initial means of communicating.'' i do find it incoherent to postulate that 2 or more people in the same time/space (or close enough to be considered ``brought together'') would have no means of communicating, completely regardless of what languages they may or may not know.

>communication is a terribly important part of what human beings
>(and all organisms!) do.

Then why is it so hard to communicate with you?

probably because we are trying to do it w/ only words. i mean that in all seriousness. i know we could get an awful lot farther, probably an awful lot more smoothly, if we were in the same room conversing w/ each other (or even if we had shared that experience sometime in the past).

cheers. --penni

Date: Mon Nov 23, 1992 7:15 pm PST
Subject: Martin's beliefs -- what for?

Martin Taylor (921123 11:40)

>When I talk about a "belief" as a structured percept, I am taking account of
>the multidimensionality of the perceptual structure. Many parallel perceptual
>signals may be components of a "belief," even if there is no actual control
>of the belief as such--no single perceptual input function to which it
>corresponds.

So you use "belief" to refer to a set of parallel perceptual signals, so $b = (p_1, p_2, \dots, p_n)$
-- belief b is a set of perceptual signals p_1, p_2 . OK so far? There may be no perceptual
input function that corresponds to this belief. I take this to mean that the set
 (p_1, p_2, \dots, p_n) is not converted into a new perceptual signal, p , where $p = f(p_1, p_2, \dots, p_n)$. Ok
still? So what is the belief for (in the model) -- what does it explain? It appears that
it might be of interest to an observer of the brain in which these signals exist as a
belief -- but of what use is it to the brain itself?

> There is only control of the multiplicity of scalar perceptual
>signals, all of whose related error signals are to be brought near zero.
>It is not necessarily satisfactory for any single function of those signals
>to be brought to a reference value, though it could be satisfactory for
>several orthogonal functions of those signals simultaneously to be brought
>to their several reference values.

Satisfactory?? What is being satisfied?? Satisfactory implies that something in the system
wants the set of signals (the belief) in a certain state. What does? How does it do it? I
really just don't understand what this concept of "belief" is designed to explain and how
it explains it.

I believe that multidimensional perceptual signals are converted functionally into scalar
perceptual variables because that is the only way I know how to model a system that
controls multidimensional variables. If you know of another way, then I would be happy to
learn it. But to me it sounds like what you are saying is that there are sets of signals
(beliefs) that are controlled as a whole -- even though they are not converted into a
scalar perceptual variable. I don't see how this works. I think you are just describing
an observable side effect of the way multidimensional control works. If 10 systems are
controlling 10 inputs you will be able to compute intercorrelations between the variables
and find significant "factors" (because, as you note, of correlations between
disturbances to high order systems that set the references for these variables). But that
is just a side effect of control -- like the fact that a person might end up writing
"Pride and Prejudice" as he resists disturbances to the knot in the rubber band demo.

But what is it that you think is explained by these "beliefs"? Just curious. They don't
seem to require any changes in the PCT model since these beliefs are not really part of
the model.

Best Rick

Date: Mon Nov 23, 1992 8:16 pm PST
Subject: Language; complex perceptions; complex control

[From Bill Powers (921123.1900)] Bruce Nevin (921123.0958) --

>When has a deaf community invented a language in the absence of
>an environing community of people possessed of language who
>intercommunicate with the deaf folks by whatever means they are able?

Gee, you guys sure do take an idea off in some strange directions. I just wanted to raise the idea of doing some experiments to see how people would actually behave in inventing a language using symbols. What's the big problem? Penni assured me this has already been done, so why not just look at the data and see how people do this? Or if by some chance the experiments were done some irrelevant way, why not just do the experiments? It might not be a very pure experiment, and it might not explore any of the more complex aspects of language, but it might provide some insights into why language is put together as it is.

>I disagree that language "is fundamentally communication". Even
>more basically, people use language to create information that
>could not exist without language and could not have existed prior
>to language.

When you get all these classifications worked out, tell me. I'm sure that there is information that couldn't exist without (verbal) language: how about "'red' is an adjective?"

>This is not to say that the perceptions ("meanings") with which
>words are correlated could not have existed, though some
>categories may well not be universal, but rather that
>higher-level organizations of those perceptions (which are
>themselves perceptions, of course) are not merely communicated by
>means of language, but are consensually created by means of
>language.

It's hard for me to see how putting words together in bunches and strings yields anything more than bunches and strings of words -- unless you've got something nonverbal to hook them up with. By putting words that are already hitched to meanings together, you can bring meanings into unforeseen mutual relationships -- but unless you're prepared to perceive the result, nothing new will happen.

>But to get at the foundations of language you must study something
>that is unique to language.

I would say you have to start with a person who has a pressing need for another person to know about or do or understand something. Language will follow. Or won't it? Gosh, how could we find out?

Martin Taylor (921123.1140) --

The editor of Theory & Psychology has rejected the paper by Tom and me, but thanks to your remarks and those of the other new reviewer, and probably to his own sense of justice, he has encouraged us to write a different kind of paper on control theory. We will probably do so.

RE: complex perceptions

>(1) as a label that we as analysts may use as a shorthand to
>communicate amongst ourselves about groups of things that may be
>seen to be happening together (in us, the patterns are the
>perceptual input functions for scalar perceptual signals, but we
>are talking about complex things happening outside of us).

It's difficult to keep the analyst and the human subject separated. It would be easier if analysts were not human, and could somehow be aware directly of "complex things happening outside of us." Unfortunately, according to PCT, that is not possible. If we are aware of something, it is a perception, even if we are analysts.

So the vexing problem arises as to how we can make models of how the brain works. The answer would seem to be "at a level below that of the functional operations of the brain." We can talk about the way computations may interact sideways at a perceptual level, so that one perception is influenced by the state of another one (Land's color theory of vision, and his experimental demonstrations, are the best example of that that I know about). But the result of these interactions is simply a set of individual perceptions. The next level has to make what it will of them, discover the interactions or not. Give me 50 more years and I may understand what I'm trying to say here.

>The next level probably makes many things of the vector, but an
>outside analyst would say that those things were of a kind
>different from the elements of the vector, or from the vector
>itself.

You may be trying to say the same thing I am trying to say. As analysts we are trying to understand the operation of a DEVICE. The outcome of this device's operation is the world that we experience.

>In a structured perceptual signal, there are many CEVs, and thus
>many degrees of freedom (or perhaps I should put it the other way
>round).

Yes, I would put it the other way around, and also add a word: in a structured perceptual signal there are many degrees of freedom, and thus many POTENTIAL CEVs. But there is no guarantee that the next level up will pick out any particular potential CEV -- $x + 1/2 y^2$ instead of $x + 1/3 y^2$.

>it might often be more efficient for the organism to have an ECS
>that directly controlled the sum,

You still miss my point. The choices for definition of a CEV are not just between "sum" and "difference." They are between $0.9a + 1.1b$, $0.95a + 1.2b$, and so on over a continuum of possible coefficients, an infinite number of possibilities. If you set the coefficients to one pair of values and control the sum, you will experience one sort of consequences. If you set the coefficients a little differently, the consequences will be a little different. So there is a scale of consequences that goes with a range of coefficients, and you have to find the place on that scale that is consistent with

controlling other variables. We find a more or less useful setting of the coefficients through reorganization. But usefulness is judged in terms of the overall error in the whole system.

>This is one version of "leaning on the world."

Yes, it is. But I try to stay away from saying what it is specifically about the world that this process leans on. It's enough for me to know that there is probably a "best" setting for the coefficients in one perceptual function, given the settings of coefficients in all the others, and given the inborn definitions of critical variables and their reference states. I'm not ready to say that the result of all this is knowledge about the world.

>Many parallel perceptual signals may be components of a "belief,"
>even if there is no actual control of the belief as such--no single
>perceptual input function to which it corresponds. There is only
>control of the multiplicity of scalar perceptual signals, all of
>whose related error signals are to be brought near zero.

When you say that these signals are "components of a belief," you produce two objects of awareness in one stroke: the components, and the belief. You are adding something to the components other than their coexistence. The moment you identify something unitary in the collection, you have gone up a level; you can now talk about degrees of a single unitary belief. If you had not introduced, even unknowingly, a higher-level perception function you would only talk about a signal, and a signal, and a signal, and a signal, as your attention moved among them. They would not be components because there would be no perception for them to be components of. It takes a higher-level perceptual function to attend to many of them at once and produce a judgment: this is a belief, and the individual signals are components of it. In fact, to make those statements (if you follow me) actually requires THREE levels: the individual signals, the belief, and the perception that the signals are related to the belief by being components of it.

It's very easy to focus on a collection of perceptions and not realize that you are perceiving them from a higher level.

>How can an ECS be developed that controls something "useful to
>control," if, at the same time, "the pattern is imposed from >above."

My point is that among the detailed signals of lower level, there are infinitely many possible functions of those signals, and any function could be controlled if its value could be affected systematically by acting on lower-level reference signals. There is no inherent "best" function of the lower signals. What makes one function more useful to control than another is in the side-effects of controlling it -- side-effects such as being able to eat or breathe. The side-effects, too, are judged relative to reference signals.

So the signals being perceived do not determine which function will ultimately be chosen. They are simply what they are. The effects of Boss Reality get into this act, of course. The organism doesn't determine the external effects of controlling one function of lower variables rather than another. It has to discover those consequences. It doesn't even determine internal effects, such as mutual exclusion of certain perceptual functions (faces vs. vases). So there are constraints on what is possible to perceive, and among those things that are possible to control, what is desirable control in terms of effects on other control systems. I think that in part you're trying to get those constraints into the argument; if so, I agree with you.

>My way of looking at it is to say that the second of these quotes
>is correct. We learn to perceive only that which we learn to
>control

I would much rather say that we learn to perceive in the process of learning to control. That is not inconsistent with your experimental evidence. You can't control something if you can't perceive it, but you can perceive many things that you can't control. However, this does not obviate the observation that what you learn to perceive is strongly influenced by the process of learning to control it. As a control system comes into being, all parts of it change shape.

>But never would I say that these patterns are "imposed from above."
>The world provides the opportunity for higher-level patterns to be
>developed out of structured perceptions, and this opportunity comes
>from below.

The world does not provide just ONE opportunity for ONE pattern to be developed. It provides a large range of possibilities (as well as impossibilities). It is not what determines in the end which possibility will be realized by a perceptual function.

>... do you get my point any differently from before? Does it matter?

It probably doesn't matter. We aren't yet at the point in our experiments and theorizing where we are forced to the only conclusion that makes current sense. I would like to get back down to the levels where there is more chance of that happening.

The next CSG meeting, to answer a private query, is indeed from July 28th through August 1, a Wednesday evening through a Sunday morning. What a treat to be able to meet you!

Oded Maler (921123) --

>consider the case when there is no explicit signal representing z ,
>but rather two low-level ECS controlling independently for x^* and
> y^* , such that $ax^*+by^*=z^*$.

In this case, we will have $x = x^*$ and $y = y^*$, because x and y are being controlled independently, per hypothesis. To be sure, there will always be a value z equal to $ax+by$, but it will not be controlled at any particular value z^* . The only way to constrain the values of x and y to satisfy $ax + by = \text{constant}$ would be to constrain x^* and y^* to fit the same condition. And to do that, a higher system would perceive a copy of x , a copy of y , and combine them to perceive $p = ax+by$. With a constant reference signal p^* we would be back to controlling along a diagonal line. Without this higher-level control, however, there is nothing to keep x^* and y^* set so they specify a point on a diagonal line.

I think you made a mistake in referring to "two low-level ECS controlling independently for x^* and y^* , such that $ax^*+by^*=z^*$." What would enforce that relationship between x^* and y^* ? If there is such a relationship, x^* and y^* are not independent.

There is a simple way to tell the difference between two systems controlling independently for x^* and y^* and a single system controlling for $z = ax+by = z^*$. In the first case, disturbing either x or y will result in direct opposition, preventing a change in the disturbed variable. The other variable will not change. In the second case, applying a disturbance to x will result in some resistance to the change in x , but will

also result in a change in y. If the output function is an integrator, so the output keeps changing as long as any error is left, the second case will result in the disturbance of x being unopposed, and all of the correction being produced by a change in y. Thank you. I will put that in the Primer series.

>... according to the current dogma in neurobiology (please correct
>me if I'm wrong) the changes in the connectivity patterns of
>neurons cease at some stage, so how are new controlled perceptual
>variables acquired?

As I now understand it, dendritic synaptic connections continue to form and break up throughout life, in human beings. No new neurons are formed, as far as I know (neurons do not divide), but both ends continue to wiggle around making new connections -- and even old synaptic connections change their weights. I think that the extent of the changes in an adult is far less than it is early in development, when neurotaxis is leading filaments toward connection sites many thousands of microns away. I'm not an expert on this subject, however. Ask a real neuroscientist.

>If I were you, I would, in parallel to the development of a full
>version for Biol. Cyb., send the rejected stuff to Nature (it's
>already in a similar format, so it costs you nothing).

Good suggestions. If the current interaction with Science doesn't produce anything, I think we will follow it. But the paper still have to be rewritten to give it a more specific motivation.

No more. Sleep.

Best to all, Bill P.

Date: Tue Nov 24, 1992 7:25 am PST
Subject: Re: foundations of language (B nevin 921123 09:58:11)

[Martin Taylor 921124 10:10]
(John Gabriel 921123 17:11 replying to Bruce Nevin (Mon 921123 09:58:11))

>>Just to head off a red herring, note that Shannon/Weaver
>>communication theory is not about communicating in this sense
>>either, any more than its near synonym information theory is
>>about meaning in any general sense of the word.

>
>Again, perhaps a dumb comment. The original Shannon Weaver paper
>was a very abstract affair. When I look at either ENIGMA (the
>biography of Turing, or the account of Bletchley Hall 1938-1945)
>or back issues of CRYPTOLOGIA, these make me inclined to argue that
>the red herring is not perhaps as much of a false trail as you might
>think.

I had intended to answer Bruce that the "red herring" is more of a "red snapper," very tasty when caught.

Bruce, I think, added that paragraph because of last year's discussion on information theory, in which I failed to get across what it is really all about. It isn't about coding and decoding, but about shifts in subjective probability distributions. I don't propose to repeat last year's postings or to paraphrase them. It's too much work for too little result. But I maintain my belief that Shannon information is at the heart of why PCT works in the real world, and why living organisms have evolved to be control systems (perhaps before long, the "to be..." will be considered a tautology).

Once a PCT interpretation of "meaning" has been developed, perhaps there might be an opportunity to investigate whether the last two lines of Bruce's paragraph have any. The Layered Protocol interpretation of "meaning" makes it quite clear that information theory is very much about "meaning in any general sense of the word."

But at present I don't think it's worth arguing about. It is only worth pursuing if we are seriously interested in getting to the foundations of PCT.

Martin

Date: Tue Nov 24, 1992 8:28 am PST
Subject: Re: Martin's beliefs -- what for?

[Martin Taylor 921124 10:40]
(Rick Marken 23 Nov 92 20:09:18 EST --that's the date in my local header)

>But what is it that you think is explained by these "beliefs"?
>Just curious. They don't seem to require any changes in the
>PCT model since these beliefs are not really part of the model.

Good, at least you've got that far. I intend no change in the model, and never did. I'm dealing with ways we talk about the model, what an external observer can see in another, and about the ways in which the control system gets itself organized. Not about a different model.

>I believe that multidimensional perceptual signals are converted
>functionally into scalar perceptual variables because that is the only
>way I know how to model a system that controls multidimensional variables.

Unfortunately, that's as recursive as the homunculus that looks at the retinal input and decides what the person sees. The set of scalar perceptual variables that the multidimensional signals are converted to is itself multidimensional, and that set quite probably could usefully be treated as a structured variable. Until you wind up with a top-level ECS that is the only one on its level, with ALL the rest of the hierarchy serving it, you can't get away from that fact. And I doubt there exists such a single high-level percept whose control entails all the rest of behaviour.

Even if it were possible to escape this homuncular recursion, it would still be valuable to talk about the set of perceptual signals that go into the perceptual input function of the (hypothetical) control-all ECS. They would form a structured perceptual signal, if only as a convenience in our discussions.

>> There is only control of the multiplicity of scalar perceptual
>>signals, all of whose related error signals are to be brought near zero.
>>It is not necessarily satisfactory for any single function of those signals
>>to be brought to a reference value, though it could be satisfactory for

>>several orthogonal functions of those signals simultaneously to be brought
>>to their several reference values.
>

>Satisfactory?? What is being satisfied?? Satisfactory implies that
>something in the system wants the set of signals (the belief) in a certain
>state. What does? How does it do it? I really just don't understand what this
>concept of "belief" is designed to explain and how it explains it.

Apparently, an unsatisfactory choice of words. What is being satisfied is the set of all the reference signals in all the ECSs. This is not a necessary consequence of the satisfaction of the reference of the single higher-level ECS that is putatively providing those reference signals as its output signals. The reverse is, however, true. If the lower set are all satisfied, then so is the single higher one. As some of us never tire of saying, PCT shows that there are many ways to kill a cat, if all you want to perceive is the dead cat. The structured set is the perceptions involved with the killing.

>So you use "belief" to refer to a set of parallel perceptual signals,
>so $b = (p_1, p_2, \dots, p_n)$ -- belief b is a set of perceptual signals p_1, p_2, \dots, p_n .
>OK so far? There may be no perceptual input function that corresponds to
>this belief. I take this to mean that the set (p_1, p_2, \dots, p_n) is not
>converted into a new perceptual signal, p , where $p = f(p_1, p_2, \dots, p_n)$.
>Ok still? So what is the belief for (in the model) -- what does it
>explain? It appears that it might be of interest to an observer of
>the brain in which these signals exist as a belief -- but of what use
>is it to the brain itself?

Are you happy if I say "none"? But the set itself is useful to the brain, and if the set were different, or if some of its members were not controlled, it would be differently (probably less) useful. It is therefore useful to us to talk about the belief, just as it is useful to chemists to talk about a molecule of "water" rather than the two hydrogen atoms bound by well understood forces to the single oxygen atom. None of the atoms know anything about this "water" (forgive the anthropomorphization). All they detect is that their probability of binding elsewhere is different from what it would be if they were off on their lonesome. It is useful in talking about the atoms to take this difference in binding probabilities for the three atoms as a property of the group, and call it "water." Shifting levels, when talking about atoms, it is useful not to have to deal with the nuclear and electronic organization within the atom. We can deal with atoms, whether or not they form "water" and regardless of their internal composition.

In a control system, it is often the case that a useful organization of a set of perceptual signals is itself the object of control through an ECS that takes the to-be-controlled pattern as its input. It is not necessary that we take this higher ECS into account when talking about the pattern. We can talk about the "water" or about the set of atoms.

As an aside, we are conscious of many perceptual signals at many levels at once, and this multidimensional consciousness cannot itself be the representation of a single scalar signal. It is something else, as yet undetermined. In everyday parlance, "belief" refers to something conscious (as does "perception") that relates to a complex of scalar perceptions. I don't use the term "belief" to refer to anything necessarily conscious, but as with "perception" I take the everyday word in a more technical context. It represents the current state of a set of perceptual signals (even if the set is as small as one member)

I find it particularly useful in the Layered Protocol theory, in which the binding to PCT depends on the multidimensional satisfaction of reference signals. It at least answers Hank Folson's question, repeated by Rick, as to what talkers and listeners are controlling for.

Martin

Date: Tue Nov 24, 1992 8:39 am PST
Subject: Random Reflections

[From John Gabriel 921124;09:02 CST]

I was thinking last night about why the net is such an interesting "place", and feel that perhaps my "field naturalist" observations might interest others. They are about things affecting communication between members, and because our ideas about "PCT of groups" are developing alongside the "PCT of individuals", they might interest others.

1. We are more diverse than the usual "community of scholars." This is possible because of INTERNET, but it's interesting because I think it's rather new to have such a large "College of All Souls at Cambridge." and so the analogue of after dinner conversation there is certainly quantitatively new, and perhaps even qualitatively new although I'm less sure about that. BBN, as the first in the field of networked discussions might reasonably say "old hat," and I would readily admit their seniority and authority.
2. As a result the average "overlap" between backgrounds is smaller than one finds in a university department, or a research group in industry.
3. This necessitates more tolerance for large differences of opinion (large weighted Hamming distances) than one finds within a single scientific discipline.
4. Large Hamming distances from referees and editors of journals come as an unpleasant surprise (I've shared Bill Powers' experience a number of times to a lesser degree and sympathise/empathise).
5. Because of 3. the temperature of flames is lower.
6. Perhaps this is an interesting experiment for all of us to watch as an exercise in building shared perception.

John Gabriel (gabriel@caesar.eid.anl.gov)

Date: Tue Nov 24, 1992 8:54 am PST
Subject: Re: Language; complex perceptions; complex control

[Martin Taylort 921124 11:10] (Bill Powers 921123.1900)

I hope that my reply this morning to Rick will also cover your posting. But if I am still missing your point, it is not the point that there are an infinite number of possible functions of any set of variables.

And

>When you say that these signals are "components of a belief," you

>produce two objects of awareness in one stroke: the components, and the
>belief. You are adding something to the components other than their
>coexistence.

In the mind of the analyst. If this happens in the mind of the one being analyzed, this is quite correct, as is all the rest of that paragraph.

Overall, in commenting on "complex perception" I don't think you have said anything about the model with which I would have disagreed, or that I would want to change. But about the implications of the model, I think there's the place where any problem lies.

Oded was quite right that if $x=y=0$, then $x+y=0$, and if there was some need for $x+y$ to be zero (bringing an intrinsic variable within tolerance bounds), controlling x and controlling y would be enough. Only if it were necessary to use variable means (such as if y were externally constrained to be non-zero) would there be any need for an ECS with a perceptual input function $x+y$.

Martin

Date: Tue Nov 24, 1992 8:58 am PST
Subject: Usenet (NetNews) Mystery

Gary

I always check both UseNet and my personal mail. Bill is definitely on UseNet at my site. In fact, everything I get from csg listserver seems to end up on UseNet.

Did you ever get my posting re: a new e-mail address for me? I never got a reply.

The old address still works so I'm sticking with it.

Regards Rick

Date: Tue Nov 24, 1992 9:58 am PST
Subject: Language exper; brain reorganization

[From Bill Powers (921124.0800)] Penni Sibun (921123.1600) --

Here's a rough undebugged version of the sort of naive experiment I have in mind. Let's make it very simple. Two people, A and B, sit across from each other at a table. There is small set of cards (20 or 30) with arbitrary symbols on them laid out on the table to one side where both can reach them. Between A and B are some objects; say a pencil, a book, a cup, and a spoon.

The object of the extremely limited game is for A and B to learn to communicate with each other about the objects on the table, and about acts related to those objects, by pointing to the cards. The rules are "as little communication as possible by other means." Ideally, none at all. But do what you have to do.

The simplest process I can think of is what I call naming, or what you say should be called referring. I'll say naming, and mean only mentally attaching a symbol on a card to an object on the table. To propose a name for one of the objects, I have to indicate an object and point to a card at the same time (pointing to a card is like uttering a word;

looking at a card is like thinking a word). You could indicate agreement by pointing to the same object and the same card, or you could disagree by pointing to the same object and a different card, or the same card and a different object.

Eventually we would reach an agreement, so when I point to a card you point to an object, or when I point to an object you point to the same card, and we are both satisfied. We now have four objects and four names. We each can draw the other's attention to a particular object by pointing to a card instead of the object.

If the task is for each person to be able to get the other to do something with the objects on the table, naming the objects is only the start. Next, I suppose, we would have to name things the objects can do, like move around, and be in relationships to other objects. To do this we would have to do things to the objects and point to cards in an effort to establish a name for what is going on. For example, I could reach out, grasp the spoon, and lift it off the table, while pointing to a SQUARE. This is an ambiguous meaning for SQUARE, because I could mean the act of lifting anything, or the state of the spoon being above the tabletop. I could lift various things, each time pointing to SQUARE, to show that it doesn't matter what I lift. Or I could point to SQUARE and the icon for one of the objects and lift the corresponding object, and do this for each object. This might indicate that SQUARE is a condition of any object, or that it is an act of lifting applied to different objects. If I indicate SQUARE alone and lift my hand with nothing in it, the meaning of SQUARE might be understood as the act of lifting.

Well, all this is just guesswork; the object is to find out what happens when real people try this. It would be informative to see how much cheating goes on -- what other means of communication have to be used to get started, or to resolve ambiguities. That would tell us how to modify the experiment and the instructions. It could turn out, for example, that the first thing people have to establish is cards meaning YES and NO.

Anyway, with very little thought behind it that is what I meant.

Oded Maler (921123) --

Thanks to Mary's sharp eye:

In Science, p. 1159, 13 Nov. 1992, there is a Technical Comment by Ramachandran, V. S., Rogers-Ramachandran, D., and Stewart, M, titled "Perceptual correlates of massive cortical reorganization." Cited is another article: T. P. Pons et. al, Science _252_, 1857 (1991), with the comment "T. P. Pons et. al. found cortical maps to be capable of an unexpectedly large degree of reorganization ... The finding extended previously recognized maximum area of cortical reorganization in adult primates from a mediolateral distance of 1 to 2 mm to about 10 to 14 mm."

Ramachandran et. al. studied phantom limb effects and found that cortical areas serving severed limbs become sensitive to stimulation from the face. Other reorganizations (in adult animals and human beings) are mentioned. In phantom limbs referred to facial areas, stimulation of the facial areas was felt as if from the limb. Interestingly, low-level (according to HPCT) signals such as warmth and pressure on the face were felt accurately, but as if they originated in the phantom limb. One-to-one mapping was preserved. I would interpret this to mean that the remapping took place at the configuration level or above, not at the sensation or intensity levels.

Tom Bourbon (921123) --

You will be getting a letter from Stam, too -- mailed at the same time as mine. It's not an unfriendly letter. Let's follow up on it. "Worlds" was an aesthetically nice piece, but I think we have to conclude that it can't be appreciated now. Let's do a basic tutorial as if to a willing audience.

Best, Bill P.

Date: Tue Nov 24, 1992 9:59 am PST
Subject: hoist by their own petard

[From: Bruce Nevin (Tue 921124 11:57:04)]

(Bill Powers (921123.1000)) --

That is indeed an excellent response to the _Science_ editor.

The reviewers assumed you were no more able to back up your claims than they would be.

Good luck! Bruce

Date: Tue Nov 24, 1992 11:04 am PST
Subject: Multidimensional perceptions & epistemology

[From Bill Powers (921124.1000)] Martin Taylor (921124.1040) --

Rick said:

>I believe that multidimensional perceptual signals are converted
>functionally into scalar perceptual variables because that is the
>only way I know how to model a system that controls multidimensional variables.

and you said:

>Unfortunately, that's as recursive as the homunculus that looks at
>the retinal input and decides what the person sees.

It's recursive if you believe that there is a pattern that is "really there" and the higher system simply recognizes it. If you believe that reorganization creates a perceptual function, which thus creates a way of combining inputs to produce a scalar output, then there's no recursion problem. The pattern does not exist in nature until a perceptual function FORMS it. After that, the brain can treat the pattern as if it has objective existence -- it can select preferred states of it, act to control it, recognize it when it occurs, see and control its relationship to other patterns, and so forth. As Wayne Hershberger said, the pattern is then "realized" -- made real. This kind of pattern-creation goes on in the mind of the analyst as well as in the analyzed brain.

I think that if we let go entirely of the idea that the perceptions in a brain have external objective correlates, we can then approach the basic epistemological problem with a fresh slant on it. I am not recommending solipsism. It is perfectly possible for an orderly universe to exist outside the bounds of experience, and to have lawful and constraining interactions with our perceptions and actions, while not resembling what we experience in the slightest. The problem of recursion arises only when we identify what we experience with what IS. When we accept that ALL perceptions are constructed, even those in the mind of the analyst, it no longer seems strange to say that a homunculus

"looks at the retinal input and decides what the person sees." Until the homunculus (read perceptual function) accepts those inputs and combines them to create a perceptual signal that depends on them in some systematic way, the person sees NOTHING. Afterwards, the person sees only a function of the inputs, basically an arbitrary function. The inputs themselves are unrepresented in the resulting perception.

There are no inherent patterns in the retinal input, save those that a human observer, using a brain, perceives there. All attempts to prove that the patterns really exist independently of perception rely on someone's perceptions. There seems to be no way out of this conclusion.

The brain's basic perceptual problem is to construct a world in which all the perceptions are mutually consistent and sufficiently controllable without excessive conflict, and in which the result of controlling the perceptions suffices to keep the organism alive and satisfying its intrinsic reference signals. Of the world that exists just outside its sensory endings, the organism knows NOTHING. It can only make a plausible model, and live within that. What I say here is simply part of that plausible model.

When you begin to see your experiences as a display, as the output of unknown transformations, you can begin to get a sense of being in a much different place from the one that's apparent. Your sense of being comfortably in charge of a familiar world is replaced by a sense of being in contact with something fundamentally beyond comprehension.

I think it was Eddington who said "The universe is not only stranger than we imagine; it is stranger than we CAN imagine." He might have been talking about vast astronomical phenomena, but he could just as well have been describing what we experience as we sit in an easy chair enjoying the warmth of a fireplace in a familiar room.

Best, Bill P.

Date: Tue Nov 24, 1992 11:18 am PST
Subject: Clinical PCT

[From Rick Marken (921124.1000)] Ed Ford (921121:0953) says:

>People are able to articulate, prioritize, and evaluate not only
>references, but set appropriate standards or rules or criteria that
>will help them reach their goals.

Well, you may be working down there where the rubber meets the road but you are dealing with some enormously precient people; apparently, they are able to know what the state of the world (disturbances) will be when they set their standards, rules and criteria so that these will be appropriate and allow them to reach their goals. How can they do this when the disturbances they will actually encounter are unpredictable and, often, undetectable? I thought that PCT made it clear that the only appropriate settings for ANY references are those that, when the outputs resulting from these reference settings are combined with prevailing disturbances, produce the intended perceptual results. Thus, you MIGHT be able to direct a persons attention of the perceptual variable that MIGHT improve their ability to CONTROL other perceptual variables (the ones that they came in complaining that they could not control) but you cannot POSSIBLY know in advance the appropriate SETTINGS for the references for these variables.

>you teach people how to use their own system, to set their own goals,

>their own standards upon which they can make choices, because PCT
>teaches me that that is how the system is designed.

This sounds right. Now the question is how one does this. I don't know how but I don't think it's by getting people to articulate "appropriate" settings for references; it must have something to do with figuring out what the person is conflicted about; conflict being the main reason for lack of control.

> The last thing I do is to impose my values or beliefs on them.

I never meant to suggest that you did; I know you don't. I am just questioning the idea (at least as you describe it and as I understand it) that one can help another person control better (which is what I imagine to be the goal of PCT therapy) by suggesting that there might be appropriate settings for ones references for any perceptual variables -- rules, standards, principles, whatever. The "appropriate" setting of a reference must vary with circumstance if the intended result is to be produced. So it's not that I think you are trying to impose YOUR values -- it is that you are suggesting that there are values that are RIGHT for the client. This is correct, as far as it goes, but the RIGHTness of that value is relative; it depends on what they are trying to achieve at a higher level (which I think you clearly understand) AND it depends on prevailing (and unpredictably changing) circumstance -- so that the setting for the value that achieved the higher order goal at one time almost certainly won't do it another time. It is this latter aspect of "setting standards" that I don't hear reflected in your ideas about PCT therapy.

Best Rick

Date: Tue Nov 24, 1992 12:05 pm PST
Subject: Re: Martin's beliefs -- what for?

[From Rick Marken (921124.1100) Martin Taylor (921124 10:40) --

I said:

>I believe that multidimensional perceptual signals are converted
>functionally into scalar perceptual variables because that is the only
>way I know how to model a system that controls multidimensional variables.

Martin replies--

>Unfortunately, that's as recursive as the homunculus that looks at the
>retinal input and decides what the person sees.

Have you still got my spreadsheet model? Have you seen any homunculi in there?

Please, help me understand what you mean here. I have heard this criticism of PCT from reviewers. Why do they think this?

>In a control system, it is often the case that a useful organization of a
>set of perceptual signals is itself the object of control through an ECS
>that takes the to-be-controlled pattern as its input. It is not necessary
>that we take this higher ECS into account when talking about the pattern.

I don't know what you mean here either. But it seems to be related to Oded's post about how systems controlling x and y independently could also be controlling x+y although there is no explicit x+y control system. I refer you to Bill's nice reply to Oded regarding the "necessity" of taking into account higher order "to be controlled" variables.

Best Rick

Date: Tue Nov 24, 1992 12:12 pm PST
Subject: Re: Multidimensional perceptions & epistemology

[Martin Taylor 921124 14:00] (Bill Powers 921124 10:00)

This is getting absurd. I continue to agree with everything you actually say about the model, and you keep assuming I am saying something that disagrees. We are clearly shooting across each other's bows, and I don't know where else to aim.

There is one point, not about the model, but about your comment, where a clarification might lead to a resolution:

>>Unfortunately, that's as recursive as the homunculus that looks at
>>the retinal input and decides what the person sees.

>

>It's recursive if you believe that there is a pattern that is "really
>there" and the higher system simply recognizes it. If you believe that
>reorganization creates a perceptual function, which thus creates a way
>of combining inputs to produce a scalar output, then there's no
>recursion problem.

The "really there-ness" of the pattern is irrelevant to the question. I do "believe that reorganization creates a perceptual function, which thus creates a way of combining inputs to produce a scalar output" but it does not relate to the recursion problem, which is in quite a different area.

The recursion is not in the development of higher-level perceptual input functions, but in that developing the higher PIFs does not get rid of the multidimensionality of the set of lower ones. If the higher ones have sufficient control that overall the lower ones are controllable by the set of higher ones, then the problem of description is simply moved one level higher. If the higher set is of lower dimensionality than the lower (and you and Rick seem to have been talking about one single higher ECS with its PIF), then there is behaviour in the lower set that is unaccounted for within the behaviour of the higher. All I am saying is that there are three reasons why it is interesting to consider the set as well as the individual, and I enumerated those reasons yesterday.

Remember that although at any one moment the reference levels for the lower set are completely determined by the outputs of the higher set, these reference levels change as the higher errors are reduced. This reduction can occur in a variety of ways, and if the dimensionality of the higher set is lower than that of the lower set, the errors can be brought to zero even if some of the lower set have their outputs blocked (by conflict or by a recalcitrant world). The higher set can be satisfied over a large (possibly infinite) set of perceptual signal levels in the lower set.

Another topic.

I'm not sure that I agree with: "The pattern does not exist in nature until a perceptual function FORMS it." It's an unrelated issue, however, and a red-herring dragged across the discussion of structured control systems.

On its own merits, the way I see it is that perceptual functions keep changing and new ones are developed under reorganization. So long as there is error (and we won't say whether the error is local or global, or in what kind of variable) there are chances for reorganization, including modifications to PIFs. When a PIF comes together with a set of output connections that result in actual control of the scalar produced by the PIF, the changes stop, or at least become less likely, less large, or less frequent. Such conditions can occur only if the PIF corresponds to some stable aspect of the world, AND if the actions that result from the output have some (probabilistically) lawful relation to the stable aspect of the world that the PIF defines.

So although a PIF can define any pattern in the world whatsoever, a stable PIF will correspond only to some aspect of the world that is both itself stable and has a stable relation to actions that the perceiver can execute. This includes the possibility of multiple relationships between actions and the aspect of the world corresponding to the PIF (i.e. what we have called the Complex Environmental Variable). Multiple means to a single end would be unnecessary if the relationship between actions and the CEV were exact and predictable, which is why the (probabilistically) was included in the previous paragraph. The more ways to skin a cat, the more probable that the cat can be skinned under a variety of circumstances. (By the way, I'm a cat-lover, in case anyone is misled by my imagery).

So, the developed structure of a mature control net is likely to contain a preponderance of PIFs that correspond to stable aspects of the world, of varying degrees of complexity. (As a corollary, it is pointless to try to model reorganization in a hierarchic network in a simulated world that does not demand as much complexity). Any control network may contain completely arbitrary PIFs, but any that do not correspond to CEVs with a reasonably lawful (probabilistic) response to action will tend to disappear. (A possible misreading of this is that superstitious behaviour will tend to disappear. That's not what it means. There's only a very slight reorganization pressure for superstition to disappear.)

Apart from those two comments, which really belong in different threads, I find nothing in your posting with which I would strongly disagree, and not much I would choose to say differently.

I wonder if it would help if you were to read my postings of the last few days once more, especially the most recent one to Rick (921124.1040).

Martin.

Date: Tue Nov 24, 1992 12:40 pm PST
Subject: Re: Usenet (NetNews) Mystery

Rick:

>I always check both UseNet and my personal mail. Bill is definiteoy on
>UseNet at my site. In fact, everything I get from csg listserver seems to
>end up on UseNet.

Good. That is how the world is supposed to me. I'll have to log into my Usenet setup again and hope that Bill appears this time.

>Did you ever get my posting re: a new e-mail address for me? I never got a >reply.

Yes I did, and I replied. And I never got an error message of undelivered mail either.

--Gary

Date: Tue Nov 24, 1992 12:42 pm PST
Subject: Re: Martin's beliefs -- what for?

[Martin Taylor 921124 14:40] (Rick Marken 921124.1100)

>>Unfortunately, that's as recursive as the homunculus that looks at the
>>retinal input and decides what the person sees.
>

>Have you still got my spreadsheet model? Have you seen any
>homunculi in there?

If I had said "as recursive as a snake eating its tail" would you have asked if I saw any snakes in the spreadsheet? Can't PCT people read similes any more? Both you and Bill seemed to read that I meant that the higher level PIF was a homunculus, and I cannot imagine how any normal reading of the quoted comparison could lead to that reading, unless it is something to which you have, unknown to me, been sensitized.

>Please, help me understand what you mean here. I have heard this
>criticism of PCT from reviewers. Why do they think this?

Since I did not provide a "criticism of PCT" of any type, the referent of "this" is unclear. If you mean that they see homunculi, I can't imagine why they think it. PCT is the most solidly behaviourist theory of psychology, the most removed from mysticism, that I have ever come across. If you mean that they say that PCT-writers use circular arguments, I haven't seen any of those, short of your argument, about which I used my simile. So I can't really answer you, without knowing what "this criticism of PCT" is.

Perhaps my other postings of today will answer your second problem.

Martin

Date: Tue Nov 24, 1992 4:28 pm PST
Subject: Blinding Flash of Light.

[From Gabriel 921124 17:29 CST]

The group theory and topology, i.e. the invariants, of the domain, which is the external world, define the range, i.e. the percepts. Just what Bill C. said to me about teaching a few days ago "Say it at the beginning, in the middle, and then again at the end!" If that is not a description of the only way for the lecture audience to experience a previously unknown invariant, I don't know what is. Saying it three different ways defines a triangulation in a 2 dimensional projection of a Kanerva space.

Also this is just a mathematician's description of gestalt theory. And the theoretical physicist's description of physical phenomena, and incidentally my formal description of semi-synonymity between computer programs. HOT DAMN!!

John Gabriel (gabriel@eid.anl.gov)

Date: Tue Nov 24, 1992 4:44 pm PST
Subject: Re (Sibun 921123) (Powers 921124)

[From Gabriel 921124; 16:21 CST]

A couple more thoughts. Penni's experiment seems to me to be able to exemplify Martin Taylor's layered protocols, and provide a shared experience for us all about that.

Also, the comment in Bill's post about reorganisation in the cortex leads me to remark further on computer programming language as language.

For me a program is a many: 1 mapping from a domain of all possible inputs (including illegal ones leading to the output ERROR!!), to a range of all possible outputs. Both the range and domain are discrete finite sets (as are human experiences if you allow for noise in the channel - continuous systems are only as real as real numbers - in fact all physical "continuous" channels have a threshold below which they don't discriminate between close pairs of inputs, if not because of Johnson noise, because of Planck's constant).

One of the things PCT must address is ?? What are the domain and range?? If they are changed environment and action, we have stimulus/response theory. If we partition the system at a cutset called "percepts", we find the traffic goes both ways across the cutset, and are led, as Bill was the first to observe, to "feedback" and "control". There's nothing wrong with even stimulus/response, provided the black box has "feedback" inside it.

I think we all agree that "percept" is an abstraction of some complex process in neuroanatomy at least for PCT of individuals, and that one day our intellectual descendants will be able to examine the detailed process in the neuroanatomy.

In the meantime, because we don't have a well defined range for percepts, it is quite easy to fall into arguments about matters we cannot decide. I think I hear some of this in what Penni has said over the last couple of months.

Now, back to my interest in programs. Here the meanings are simpler, but the range is immediately accessible. So that for example I can use topology and group representations in study of programs, but any attempt to do the same for PCT is more like the Chicago school of thought about continuous groups in embryology 30-40 years ago than it is like PCT.

The central issue about programs at present is the analogue of story understanding in language, and what I have to say was prompted by the discussion of synonyms in language.

If one rewrites a program in one language into another, say from FORTRAN into "C", or more relevant to today's problems, from a VAX sequential architecture, to an INTEL Touchstone parallel architecture, one would like it to "do the same thing", i.e. implement the same mapping from domain to range, or "tell the same story" in a very

special sense, since the mapping from domain to range is a global property of the program, just as story told by a book is a global property of the book.

For the purposes of the user, two programs, one in "C" for the Touchstone, and one in FORTRAN for the VAX, that implement the same mapping are synonymous except for the compilers and job control language that they need to make them run.

There is a completely unambiguous statement of this in terms of the various kinds of mappings involved - "C" to Touchstone executable, FORTRAN to VAX executable, Touchstone to VAX input file, and VAX output to Touchstone output.

Some of these mappings - the ones involving program or compiler execution are in a way "internal" to the agents doing the machine understanding - the computers and their operating systems, others are "external" - the input/input mappings and the output/output mappings.

And perhaps the process of iterative refinement of the new program to make it do that same thing as the old one did, is a PCT like control process by the programmer on the code for the new program. Like a set of "layered protocols" to "tell the right story" to the Touchstone.

BY GOSH - I managed by the skin of my teeth to say "control" and thus qualify this post for CSGNET. Would I have otherwise had to publish it elsewhere?? - Bill P. don't be offended, this is a very gentle legpull, and I have an enormous respect and admiration for what you have done.

John Gabriel (gabriel@eid.anl.gov)

Date: Tue Nov 24, 1992 5:19 pm PST
Subject: Science letters, late

HI Tom, Greg, and Rick --

The following letters and the implied materials will be in the mail today. I gave myself two days to feel sorry for Greg and myself. Most of this was Mary's idea -- I don't know why I didn't think of referring to the Technical Comments reviews when we first submitted the article. It probably would have made a difference. Anyway, let's see what comes of all this.

Bill.

November 23, 1992

Editors
Science

Dear Sirs,

I'm enclosing two copies of a computer model of human pointing behavior in two stamped but unaddressed envelopes. I am asking you to forward them to two anonymous viewers of a Technical Comment submitted by me (and rejected) about a year ago (Ref. TC 34076). I have the address of the third reviewer, Bizzi, and am mailing the material directly to him.

Behind this request is a short story.

I recently submitted an article (Ref. RE924541) with Gregory Williams on human pointing behavior to Science. This article has also just been rejected. The main motive behind submitting this article came from comments of the three reviewers of the rejected technical comment, of which reviews excerpts follow.

Reviewer 1 said

"Although I agree with the general point of his [my] comment, I feel that it represents mainly a misunderstanding of ... Bizzi. I would very much like to see the development of a more general model, one that incorporates the basic observations of Bizzi et. al. and the philosophy of Powers ... but I would rather see the Bizzi et. al. point of view countered with data and research."

Reviewer 2 said

"Finally, the main point of Power's [sic] technical comment seems to be an advertisement for his model ("Soon to appear in a journal in your neighborhood.") Certainly, if his model can do all he claims, it would be of interest, but too much of this technical comment is based on acceptance of Power's stipulation that his model, offered for mid-1992 delivery, works. However, in the absence of detail, his comments fall flat."

Emilio Bizzi (the target of the technical comment) said

"In summary, we believe the burden of proof is on Dr. Powers to show that it is indeed simple for the CNS to generate posture and movement using a feedback control scheme which resides predominantly in the CNS and which must combine the diverse sensor information sources he describes. Such a scheme must surely have to take into account the various different feedback loop delays, and the complexity of muscle, sensor and limb dynamics."

The article recently rejected was intended to supply the missing information about the model and its performance. In the interests of brevity and to avoid politics, we made no reference to the earlier rejected technical comment, but offered a compact discussion of the model to fit the space restrictions of a research article. We also included discs with a runnable version of the model and a more extensive writeup to make its operation clear.

Unfortunately, the article was apparently submitted at the same time as other articles of superior interest, relevance, novelty, and so on, and failed to pass the initial cut by the board of reviewers. As a result, the reply motivated by the three reviewers of last year will not appear in Science. That is why I ask you to forward these copies of the program to those reviewers, so at least they can see that my comments were not based on thin air.

Thank you for your help.

William T. Powers
73 Ridge Place, CR 510
Durango, CO 81301

forwarded by Science

November 23, 1992

Dear Sir,

In September 1991, you wrote one of two anonymous reviews rejecting a technical comment by me on an article by Bizzi. Your comment was one of the two following:

Reviewer 1 said

"Although I agree with the general point of his [my] comment, I feel that it represents mainly a misunderstanding of ... Bizzi. I would very much like to see the development of a more general model, one that incorporates the basic observations of Bizzi et. al. and the philosophy of Powers ... but I would rather see the Bizzi et. al. point of view countered with data and research."

Reviewer 2 said

"Finally, the main point of Power's [sic] technical comment seems to be an advertisement for his model ("Soon to appear in a journal in your neighborhood.") Certainly, if his model can do all he claims, it would be of interest, but too much of this technical comment is based on acceptance of Power's stipulation that his model, offered for mid-1992 delivery, works. However, in the absence of detail, his comments fall flat."

I recently submitted a research article on the model in question to Science. Unfortunately it was rejected, so I lost my opportunity to provide the missing information about the model. I have asked Science to forward this collection of materials to you (as I don't know who you are), so that at least you will see that the model does exist and does perform as claimed.

Enclosed are (1) the rejected article, (2) a writeup of the "Little Man" model, and (3) a 360-K computer disk containing the model in runnable form, along with its source code. I advise running the model on a fast AT, 386, or 486 computer. On a 486-33Mhz computer it runs essentially in real time.

Yours truly,

William T. Powers
73 Ridge Place, CR 510
Durango, CO 81301
(303) 247-7986

November 23, 1992

Dear Dr. Bizzi,

In September 1991, you wrote one of three reviews rejecting a technical comment by me on an article of yours. The relevant part of your comment was:

"In summary, we believe the burden of proof is on Dr. [Note: Mr.] Powers to show that it is indeed simple for the CNS to generate posture and movement using a feedback control scheme which resides predominantly in the CNS and which must combine the diverse sensor information sources he describes. Such a scheme must surely have to take into account the various different feedback loop delays, and the complexity of muscle, sensor and limb dynamics."

The other two reviewers said essentially the same thing: show the model.

I recently submitted a research article on the model in question to Science. Unfortunately it was rejected, so I lost my opportunity to provide the missing

information about the model. I am sending this collection of materials to you, so that at least you will see that the model does exist and does perform as claimed. I have asked Science to forward a similar package to the two anonymous reviewers.

Enclosed are (1) the rejected article, (2) a writeup of the "Little Man" model, and (3) a 360-K computer disk containing the model in runnable form, along with its source code. I advise running the model on a fast AT, 386, or 486 computer. On a 486-33Mhz computer it runs essentially in real time.

Yours truly,

William T. Powers
73 Ridge Place, CR 510
Durango, CO 81301
(303) 247-7986

Date: Tue Nov 24, 1992 6:45 pm PST
Subject: Re: Language exper

(penni sibun 921124.1700)

[From Bill Powers (921124.0800)]

Penni Sibun (921123.1600) --

Here's a rough undebugged version of the sort of naive experiment I have in mind. Let's make it very simple. Two people, A and B, sit across from each other at a table. There is small set of cards (20 or 30) with arbitrary symbols on them laid out on the table to one side where both can reach them. Between A and B are some objects; say a pencil, a book, a cup, and a spoon.

The object of the extremely limited game is for A and B to learn to communicate with each other about the objects on the table, and about acts related to those objects, by pointing to the cards. The rules are "as little communication as possible by other means." Ideally, none at all. But do what you have to do.

i think your idea is interesting, and certainly simple enough to be doable. my prediction w/b that many subjects would find the task difficult to understand, frustrating, and difficult to carry out, to various degrees, but all in fascinating ways. i think it would be an excellent source of hypotheses for what sorts of communication get used.

following are some examples of ``referring'' as opposed to ``naming'': pointing at; nodding towards; gesturing vaguely; rolling one's eyes; saying this, that, you know, the thingamybob, something, it, him, her, them,....

notice that none of these involves a name. i don't know of any actual statistics, but clearly some significant fraction of references are nameless.

cheers. --penni

Date: Tue Nov 24, 1992 9:59 pm PST
Subject: Error-free programming

[From Dag Forssell (921124)]

I would very much like to share the idea of error-free, control based programming with a PhD candidate in Computer Science, who will be visiting thanksgiving. (Daughter's boyfriend; he is about to select subject for thesis project). I seem to recall at least two postings on the subject last spring. Searching "program," beginning Nov 91, I have found one of the references:

>Bill Powers (920319.0900) >Keynote address:
>
>Possibly another subject of interest might be writing error-free
>programs. I've always thought that there isn't much distance between
>current programming practices and an HPCT approach. Instead of treating
>the computer as an open-loop device, monitor every intermediate result
>and compare it with a reference signal to make sure it's of the right
>kind, makes sense in some terms, and so on.....

Can someone with a good memory or a good file system help me find the other(s)?

Thanks, Dag

Date: Wed Nov 25, 1992 6:45 am PST
Subject: Re: Error-free programming

[From gabriel 921125; 07:41 CST] >[Ref Dag Forssell (921124)]

>I would very much like to share the idea of error-free, control based
>programming with a PhD candidate in Computer Science, who will be
>visiting thanksgiving. (Daughter's boyfriend; he is about to select
>subject for thesis project). I seem to recall at least two postings on
>the subject last spring. Searching "program," beginning Nov 91, I have
>found one of the references:

>>Bill Powers (920319.0900) >Keynote address:
>>
>>Possibly another subject of interest might be writing error-free
>>programs. I've always thought that there isn't much distance between
>>current programming practices and an HPCT approach. Instead of treating
>>the computer as an open-loop device, monitor every intermediate result
>>and compare it with a reference signal to make sure it's of the right
>>kind, makes sense in some terms, and so on.....

>Can someone with a good memory or a good file system help me find the other(s)?

>Thanks, Dag

I think you may be referring to what the software engineering community call regression analysis. Mail to bkero@eid.anl.gov for more info.

Another point of view is put forward by R.S. Boyer & J. Strother Moore "The Correctness problem in Computer Science", Academic Press 1981, and two other books on Boyer & Moore's

theorem prover, one published about 1979, and the other about 1989. Email to Bob Boyer at UTEX Austin for more - I'll try to dig out the good E-mail address later. But Bob is VERY mathematical and rigorous, and you may be uncomfortable with him.

Yet another view is held by colleagues at Argonne who write resolution based theorem proving systems. See the book by Wos and others on Automated Reasoning about 7 years ago or mail to lusk@mcs.anl.gov, wos@mcs.anl.gov, overbeek@mcs.anl.gov or chisholm@eid.anl.gov. Greg Chisholm is a very friendly chap happy to talk. I may not have the E-mail address quite right but I can correct it later if need be.

Another approach to the problem is embedded in Bill Wulf's specification language IDL, see the book by Richard Snodgrass et alia "The Interface Definition Language, Computer Science Press 1988.

There is also a very fine resolution based school at SRI, the book and dramatis personae escape my recollection at the moment.

The world's best group in many ways is at the new Institute for Mechanical Reasoning founded a couple of years ago by the Beijing Academy of Sciences. Wu Wen Tsun, who gained an international reputation for algebraic topology before winning a second one for automated theorem proving leads that effort when he can find time away from the serious business of advising his Govt. I will dig out his E-mail address and send it later.

Finally, my own methods are based on G.W. Mackey's Theory of Induced Group Representations which is now about 35 years old. I sent a paper off to Prof Wu a while back for an international conference in China, but was not able to get away from here to accept my invitation to attend, so I think Wu may be a little hurt. I'm not sure if it has been published. I am working sporadically on a successor, but my views are about as far from the US norm as PCT is from its US norm.

There are also active schools in Europe, particularly in the UK and in Germany, the universities of Karlsruhe and Edinburgh come to mind.

There is a parallel effort in proof of correctness for VLSI which has substantial commercial value. There is a very good group at Cornell, and others in places like IBM, Edinburgh, XEROX (Harry Barrow at PARC comes to mind, perhaps Penni can find him for you). I can put you in touch with a student of mine presently at Cornell once I dig thru my files.

John gabriel (gabriel@eid.anl.gov)

Date: Wed Nov 25, 1992 9:39 am PST
Subject: Re: Error-free programming

[Hans Blom 921125]

A name that comes to mind is Edsger "GOTO considered harmful" Dijkstra, formerly employed at the place where I am employed, now somewhere in Texas (Austin?). He always harps on the impossibility of correctness proofs (NP-complete), but shows an alternative: a design methodology. Pre- and post- conditions surround the code in question and assert what must be true at that point in the code. Thus the code is a transformation of variables, where the input to the code is checked (e.g. if the code calculates a square root of x , x must be positive) as well as the output (if y is the square root of x , $y * y$ must be equal to x). The hope is, of course, that program modules that are known to be

locally correct will be reliable building blocks in a larger program. Modularization and information hiding are the key words here. In PCT-terminology: at the concept level, you want to be concerned about the level immediately below only.

Hans Blom

Date: Wed Nov 25, 1992 2:35 pm PST
Subject: Re: Usenet (NetNews) Mystery

To: g-cziko@UIUC.EDU
Orig-To: "Gary A. Cziko" <g-cziko@UIUC.EDU>
Subject: Re: Usenet (NetNews) Mystery
Newsgroups: bit.listserv.csg-1
References: <199211241955.AA27900@ux1.cso.uiuc.edu>

Hello, Gary -

Speaking of Netnews, did you see the new group that just came out? It's called "alt.control-theory." The creator was getting soundly flamed for creating a "useless" group, citing sci.engr.control as the existing group for that topic. Did a CSGer create that group? I believe the creator's name is Richard Banks (rbanks@nyx.cs.du.edu). Just thought you might find that interesting.

Denny Powers
c/o dormont@spot.colorado.edu
or dennyp@exabyte.com

Date: Wed Nov 25, 1992 2:41 pm PST
Subject: Re: Usenet (NetNews) Mystery

[from Gary Cziko 921125.2100 GMT]
Denny Powers (921125) says:

>Speaking of Netnews, did you see the new group that just came out?
>It's called "alt.control-theory." The creator was getting soundly
>flamed for creating a "useless" group, citing sci.engr.control as
>the existing group for that topic. Did a CSGer create that group?
>I believe the creator's name is Richard Banks (rbanks@nyx.cs.du.edu).
>Just thought you might find that interesting.

No, I know nothing about this new group or who Richard Banks is. I'll have to check into this.

By the way, does dear old dad show up for you on bit.listserv.csg-?

Date: Wed Nov 25, 1992 2:41 pm PST
Subject: Re: Martin's beliefs -- what for?

[From Oded Maler (921125)

* Rick Marken (921124.1100):
*
*

* I don't know what you mean here either. But it seems to be related to
* Oded's post about how systems controlling x and y independently
* could also be controlling x+y although there is no explicit x+y
* control system. I refer you to Bill's nice reply to Oded regarding the
* "necessity" of taking into account higher order "to be controlled"
* variables.

Although in general relying on Bill's answers is a good practice, I think in this case it is not so. I think I demonstrated how a system can control independently x and y, and as a trivial side-effect control any combination of them. If one such combination is all what an observer sees, he will conclude that this combination is a controlled perceptual variable. There are probably more complicated situations in which having a dedicated signal for a combination is better - and I would like to hear which.

Anyway I repeat my feeling that this issue might shed a new light on the question of "real" vs. "as-if" in the analysis of purposeful behavior. It might be also related to phenomena such as the following: when we learn a new task (e.g., skiing, to take something I'll have to do soon) it is well known that after some time this control "know-how" migrate downwards from the conscious to the unconscious (sp.?) and maybe this is related to the transition from not having to having a dedicated signal.

I don't pretend to have invested a lot of thought and work on reorganization as others on this net had, so I apologize if these anecdotal speculations do not make sense.

(Bill Powers (921124.0800))

(About neurological reorganization in adults)

It's a relief. Otherwise if we accept "signal for every percept" the only way to learn a new thing is to forget another...

--Oded

Date: Wed Nov 25, 1992 3:40 pm PST
From: mmt
EMS: INTERNET / MCI ID: 376-5414
MBX: mmt@ben.dciem.dnd.ca

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Re: Error-free programming

Dag,
I find no occurrences of "error-free" since January, except in reference to language phenomena, except the Powers posting you found.

Martin

Date: Wed Nov 25, 1992 3:43 pm PST
Subject: Re: Martin's beliefs -- what for?

[From Rick Marken (921125.1430)]

OK. I think I'm getting the hang of it here.

Let's try a real live reply to

Oded Maler (921125)

who says:

>Although in general relying on Bill's answers is a good practice, I
>think in this case it is not so. I think I demonstrated how a system
>can control independently x and y, and as a trivial side-effect
>control any combination of them.

Maybe you demonstrated how independent control of x and y can "appear" to be control of any combination of x,y; but anyone who understood what the term "control" meant would not take appearances for granted; s/he would do the test for the controlled variable and find that f(x,y) is not controlled at all. I think you have demonstrated nothing more than the fact that the behavior of any variable or function of variables can look like control (ie. a purposefully produced result); and that is one of the misconceptions about observed behavior that PCT is trying (with little apparent success) to correct.

>Anyway I repeat my feeling that this issue might shed a new light on
>the question of "real" vs. "as-if" in the analysis of purposeful behavior.

Yes. It certainly does. The test for the controlled variable would quickly and simply reveal that x is controlled, y is controlled and f(x,y) is not controlled.

In fact, I have done a demo of exactly this sort --in a version of the mindreading program that can tell if the subject is controlling x, y and/or x-y; it works like a charm. Bill P. described a similar demo in his Psych Review paper.

Best regards Rick

Date: Wed Nov 25, 1992 4:10 pm PST
Subject: Misc subjects

[From Bill Powers (921125.1400)]

Martin Taylor (921125) --

>This is getting absurd. I continue to agree with everything you
>actually say about the model, and you keep assuming I am saying
>something that disagrees. We are clearly shooting across each
>other's bows, and I don't know where else to aim.

Communication is like playing Battleship. Shots in the dark. Might as well just keep feeling our way and hoping for a hit.

>>It's recursive if you believe that there is a pattern that is
>>"really there" and the higher system simply recognizes it.

>The "really there-ness" of the pattern is irrelevant to the
>question...The recursion is not in the development of higher-level
>perceptual input functions, but in that developing the higher PIFs

>does not get rid of the multidimensionality of the set of lower
>ones.

If there is a set of lower-level controlled perceptions, they will be in the states dictated independently by their respective reference signals, won't they? So their "pattern" will reflect only the pattern imposed by reference signals from higher systems, not any real pattern in the environment.

I echo Rick's question: what difference does it make that there are multiple variables if nothing about that fact affects the rest of the system? To which I add, if there is a fact about the system that does not relate to relationships among variables in the system, why should the analyst bother with it?

>If the higher [systems] have sufficient control that overall the
>lower ones are controllable by the set of higher ones, then the
>problem of description is simply moved one level higher.

I don't agree. In my model, if not in yours, the different levels control for logically different types of variables. When you go up a level you're describing something new -- transitions, for example, instead of configurations. A new kind of perceptual operation has to be involved, in this case computing derivatives. You can't build the hierarchy out of identical computing elements at every level. At least not my hierarchy.

>Remember that although at any one moment the reference levels for the
>lower set are completely determined by the outputs of the higher set,
>these reference levels change as the higher errors are reduced. This
>reduction can occur in a variety of ways, and if the dimensionality of
>the higher set is lower than that of the lower set, the errors can be
>brought to zero even if some of the lower set have their outputs
>blocked (by conflict or by a recalcitrant world). The higher set can
>be satisfied over a large (possibly infinite) set of perceptual signal
>levels in the lower set.

Yes, I think I was saying the same thing. To me, this means that there is no necessary set of patterns in the lower set, because ANY orthogonal set of higher-level functions would allow independent control of the higher perceptions.

>So although a PIF can define any pattern in the world whatsoever, a
>stable PIF will correspond only to some aspect of the world that is
>both itself stable and has a stable relation to actions that the
>perceiver can execute.

This comes back to our very different conceptions of reorganization. I say that there are intrinsic variables that are NOT directly involved in the operation of the hierarchy but are affected by its operations. You have proposed that the intrinsic variables are aspects of the hierarchy itself. The present proposal is an example of that. If the only criterion for choosing a perceptual organization is some principle of stability or least action, then it is the physics of the world and the nervous system that determine what we perceive, not the state of the organism. This leaves open the question of why satisfying some physical principle of stability of perception should have any beneficial effect for the organism.

I think that there are multitudes of alternative ways of perceiving the environment that are essentially equivalent with respect to controllability. I think that the choice

between them is not dictated by satisfying some principle, but by the requirement that choosing certain of the controllable perceptions turns out to help the organism survive. I do not learn to control my food with a knife and fork because that mode of control satisfies a stability condition, but because it is one way to get food inside me that also satisfies the arbitrary demands of others around me.

I think you are looking for a way to define some condition inside the hierarchy itself such that satisfying it will result in a viable control hierarchy. I don't think that is the right place to look, in general, although there are certain content-free aspects of the hierarchy such as error signals that do have significance for survival. I think a more general and more likely model must ask what it is about the organism that could be monitored that would indicate its state of functioning. That does not have to be the same as the system that gets reorganized. And that kind of model gives a plausible evolutionary reason for reorganizing; I don't find abstract principles to be very plausible ways of making a system actually work.

>So, the developed structure of a mature control net is likely to
>contain a preponderance of PIFs that correspond to stable aspects
>of the world, of varying degrees of complexity.

I don't disagree with this; I only add that these are not the only stable aspects of the world (actually, of the set of perceptual functions) that could be chosen. What is stable, by the way, is not just the world, but the world-as-perceived. A stable perceived world does not necessarily reflect naturally stable aspects of the world. You can perceive a balanced broomstick in such a way (including necessary derivatives) that you can control it in an naturally unstable state. Your negative feedback overrides the broom's natural positive feedback.

John Gabriel (921124.1525) --

>For me a program is a many: 1 mapping from a domain of all possible
>inputs (including illegal ones leading to the output ERROR!!), to a
>range of all possible outputs.

I suppose my definition is equivalent, but the way I've put it for purposes of defining a program level in the model is that a program is a network of choice-points. Some of the "inputs" at least do not have to be present-time inputs; they can even be built into the system, as when you say "if $x + y > 4$ " then ---. The "4" is a built-in constant, put in when the program was made.

In my model I split out the sequence level as separate, because sequences can be controlled without any choice-points (recipes). So a pure program consists of inputs, tests or comparisons, and outputs, where the inputs and outputs are sequences (including the degenerate case of a sequence of one element).

>Both the range and domain are discrete finite sets (as are human
>experiences if you allow for noise in the channel - continuous
>systems are only as real as real numbers - in fact all physical
>"continuous" channels have a threshold below which they don't
>discriminate between close pairs of inputs, if not because of
>Johnson noise, because of Planck's constant).

That's pushing it a bit far. You're referring everything to digital computers, whereas the nervous system is probably an analogue computer except at levels where the analog elements are kludged up to do digital operations (as in the case of electronic "digital"

computers built with analog transistors). The natural unit of neural signals is frequency, which correlates with average concentrations of neurotransmitters and messengers inside cell bodies and degree of tension in muscles. Planck's constant is a long, long way below the level of description that matters even in the behavior of a single neuron.

Also, the existence of noise in a channel is not the same as quantization of the signal-carrying variable. There is no discrete "threshold" of detection in a continuous channel, either for nonzero signals or for differences between signals. There are continuous curves of discriminability, or as Martin would say, probability of discriminating. As in any noisy channel, you can trade time against discrimination; the longer you want to take to average the signal, the better the discrimination gets. This is not true, by the way, in a digital computer: the number of bits in a number is the best you can do (without cheating by emulating longer word-lengths).

>One of the things PCT must address is ?? What are the domain and range??

I'll let you address it. I'm more interested in getting models to work like the real system.

>So that for example I can use topology and group representations in
>study of programs, but any attempt to do the same for PCT is more
>like the Chicago school of thought about continuous groups in
>embryology 30-40 years ago than it is like PCT.

Oh, THAT Chicago school of thought. Huh?

>For the purposes of the user, two programs, one in "C" for the
>Touchstone, and one in FORTRAN for the VAX, that implement the same
>mapping are synonymous except for the compilers and job control
>language that they need to make them run.

The concept of "mapping" sounds static to me. In programming control systems on digital computers, it isn't enough to get the logic and the connectivity right; you also have to handle physical time correctly. If you program a control system with variable loop gain on a serial machine, but don't take time properly into account, you'll conclude that the maximum negative loop gain must be less than 1. In real control systems, loop gains can be in the millions (negative). This problem doesn't have to be handled in a parallel machine because things that are actually happening at the same time can run at the same time. But you still need scaling to represent physical time -- that is, how much physical time is represented by one iteration. That determines how much a mass will accelerate under a given force, and so on.

Penni sibun (921124.1700) --

>my prediction w/b that many subjects would find the task
>difficult to understand, frustrating, and difficult to carry out,
>to various degrees, but all in fascinating ways.

My idea, too. Finding out just what is difficult about it should give us some basic information.

Roger on the examples of referring. I guess referring is not what I had thought of as naming.

Best, Bill P.
Date: Wed Nov 25, 1992 7:15 pm PST
Subject: single vs multiple control: no problem to distinguish

[From Bill Powers (921125.2000)]

Oded Maler (921125) --

(written before seeing Rick's reply)

>I think I demonstrated how a system can control independently x and >y, and as a trivial side-effect control any combination of them. >If one such combination is all what an observer sees, he will >conclude that this combination is a controlled perceptual variable.

If x and y are independently under control, then the only way to make sure that x and y are related by $ax + by = \text{constant}$ is for something or someone to vary the reference signals x^* and y^* , and monitor for the condition $ax^* + by^* = \text{constant}$. Of course if you trust the control systems to make $x = x^*$ and $y = y^*$ at all times, then all you need to do is program x^* and y^* so that $ax^* + by^* = \text{constant}$. Then, as you say, it will seem that there's a control system for the function $ax + by$. In actuality, this would be an open-loop system at the level that's setting the reference signals.

However, if the observer is allowed to apply disturbances to either x or y, it can quickly be discovered that this is not really a control system. If $ax + by$ were actually being sensed and controlled directly, a disturbance of x would produce a compensating change in y and vice versa. This will not happen if the two control systems are independent and the reference signals are related by $ax^* + by^* = \text{constant}$. The x control system alone will react to a disturbance of x, and the y system alone to a disturbance of y.

>Anyway I repeat my feeling that this issue might shed a new light
>on the question of "real" vs. "as-if" in the analysis of purposeful behavior.

In the scientificphilosophical literature, questions like this are often posed and answered as if the answer somehow had to be derived through passive observation of appearances -- a hands-off approach. But there is also a hands-on approach that involves doing something to the system in the attempt to distinguish between possible explanations: the experimental approach. I favor the experimental approach, because the hands-off or passive observation approach is not good at resolving ambiguities.

>It's a relief. Otherwise if we accept "signal for every percept"
>the only way to learn a new thing is to forget another...

What, you mean that some people can learn something new without forgetting two other things?

Best, Bill P.

Date: Wed Nov 25, 1992 8:20 pm PST
Subject: Re: Misc subjects

[Martin Taylor 921125 22:50] (Bill Powers 921125.1400)

Happy Thanksgiving, USA.

I think I'll give the discussion of why it is useful to think and talk about sets of ECSs a pass for a while, until a new way of approaching it occurs to one of us. If once more it turns out to be useful in a particular instance, maybe that will give a lead. Meanwhile, let me just once more reiterate that I have no disagreement with you on the workings of a mature hierarchy, or that scalar perceptual signals are what are controlled (en masse).

>>If the higher [systems] have sufficient control that overall the
>>lower ones are controllable by the set of higher ones, then the
>>problem of description is simply moved one level higher.

>

>I don't agree. In my model, if not in yours, the different levels
>control for logically different types of variables. When you go up a
>level you're describing something new -- transitions, for example,
>instead of configurations. A new kind of perceptual operation has to be
>involved, in this case computing derivatives. You can't build the
>hierarchy out of identical computing elements at every level. At least
>not my hierarchy.

As far as I can see, your model does use the same kind of computing elements at every level. All that changes is the perceptual input function, isn't it? A configuration-level function has no delays, whereas a sequence-level PIF is basically a shift register, and a category-level PIF involves a catastrophe. If that's what you mean by non-identical computing elements, then I have no problem. To me, they are identical, in that they can all be described as multiple inputs to a PIF that produces a scalar value that is compared with a reference scalar value and the error transformed to a signal that goes down to lower ECSs. When I have tried to lay this out in detail, to ensure that we have been talking about the same thing, you have always agreed with this description. I think it is yours, isn't it?

On reorganization, there isn't so much a disagreement as that I don't see how your version of reorganization can be guaranteed to be the only one. I do admit to assuming from time to time my "integrated local" model when I talk about reorganization, but I don't think it matters in most arguments. Let's see if it does in this instance.

>>So although a PIF can define any pattern in the world whatsoever, a
>>stable PIF will correspond only to some aspect of the world that is
>>both itself stable and has a stable relation to actions that the
>>perceiver can execute.

>

>This comes back to our very different conceptions of reorganization. I
>say that there are intrinsic variables that are NOT directly involved in
>the operation of the hierarchy but are affected by its operations. You
>have proposed that the intrinsic variables are aspects of the hierarchy
>itself. The present proposal is an example of that. If the only
>criterion for choosing a perceptual organization is some principle of
>stability or least action, then it is the physics of the world and the
>nervous system that determine what we perceive, not the state of the
>organism. This leaves open the question of why satisfying some physical
>principle of stability of perception should have any beneficial effect
>for the organism.

>

>I think that there are multitudes of alternative ways of perceiving the

>environment that are essentially equivalent with respect to
>controllability. I think that the choice between them is not dictated by
>satisfying some principle, but by the requirement that choosing certain
>of the controllable perceptions turns out to help the organism survive.
>I do not learn to control my food with a knife and fork because that
>mode of control satisfies a stability condition, but because it is one
>way to get food inside me that also satisfies the arbitrary demands of
>others around me.

Your para 1 of this comment: It isn't so much that the physics of the world and the nervous system determine what we can perceive (though this is true), but that they determine what percepts we can control, and therefore what percepts can be useful to the organism. Uncontrollable percepts seem to me to be a waste of resources, and thus something that evolution would probably not allow us much of. Stability of perception is not quite the right term, here (even if I did use it, which I don't recall). What I am after is the stable PIF, which belongs in an ECS that can do some good because it can control.

This isn't to say that all controllable percepts are or can be useful (in maintaining intrinsic variables near their reference values). It is to say that uncontrollable percepts are not useful, and therefore are unlikely to come from stable PIFs.

Your para 2: I agree entirely that there are multitudes of alternative ways...that are essentially equivalent with respect to controllability. I had intended to go into that point, but I thought I was already posting too many words. But since you bring it up--even within the same set of dimensions (the same set of lower-level perceptual signals, for example), there are an infinite number of possible sets of orthogonal representations at higher levels, especially when (as you point out) there are many different kinds of functions used in these representations. When the system is restricted only by the world, the dimensionality itself is indeterminate. But none of this argues against the probability that there are an even greater number of unstable ways that the world could, in principle, be perceived. It's rather like "there are more real numbers than integers, but we count only in integers."

Since PIFs that lead to uncontrollable percepts will not help bring intrinsic variables near their reference values under your reorganization scheme or mine, I don't think our difference there matters to the argument.

Martin

Date: Wed Nov 25, 1992 8:42 pm PST
Subject: J. Adaptive Behaviour (from GA-LIST)

[Martin Taylor 921125 23:30]

Here are the contents pages from the first two issues of J. Adaptive Behavior, copied from the GA-LIST digest, 6-38.

=====

From: meyer@biologie.ens.fr (Jean-Arcady MEYER)
Date: Tue, 3 Nov 92 16:17:27 +0100
Subject: Adaptive Behavior Journal - Table of Contents

The first issue of Adaptive Behavior was released in August 1992. The second

is under press.

For inquiries or paper submissions, please contact one of the editors:

- Editor-in-Chief: Jean-Arcady Meyer, France - meyer@wotan.ens.fr
- Associate Editors: Randall Beer, USA - beer@alpha.ces.cwru.edu
Lashon Booker, USA - booker@starbase.mitre.org
Jean-Louis Deneubourg, Belgium - sgoss@ulb.ac.be
Janet Halperin, Canada - janh@zoo.utoronto.ca
Pattie Maes, USA - pattie@media-lab.media.mit.edu
Herbert Roitblat, USA - roitblat@uhunix.uhcc.hawaii.edu
Ronald Williams, USA - rjw@corwin.ccs.northeastern.edu
Stewart Wilson, USA - wilson@smith.rowland.com

=====

ADAPTIVE BEHAVIOR 1:1
Table of Contents

A Model of Primate Visual-Motor Conditional Learning
by Andrew H. Fagg and Michael A. Arbib

Postponed Conditioning: Testing a Hypothesis about Synaptic Strengthening
by J. R. P. Halperin and D. W. Dunham

The Evolution of Strategies for Multi-agent Environments
By John J. Grefenstette

Evolving Dynamical Neural Networks for Adaptive Behavior
By Randall D. Beer and John C. Gallagher

=====

ADAPTIVE BEHAVIOR 1:2
Table of Contents

Adapted and Adaptive Properties in Neural Networks Responsible for Visual
Pattern Discrimination.
By J.-P. Ewert, T.W. Beneke, H. Buxbaum-Conradi, A. Dinges, S. Fingerling,
M. Glagow, E. Schurg-Pfeiffer and W.W. Schwippert.

Kinematic Model of a Stick Insect as an Example of a 6-legged Walking System.
By U. Muller-Wilm, J. Dean, H. Cruse, H.J. Weidemann, J. Eltze and
F. Pfeiffer.

Evolution of Food Foraging Strategies for the Caribbean Anolis Lizard using
Genetic Programming.
By J.R. Koza, J.P. Rice and J. Roughgarden

Behavior-based Robot Navigation for Extended Domains.
By R.C. Arkin

=====

Martin

Date: Wed Nov 25, 1992 8:49 pm PST
Subject: Re: Misc subjects

Bill, let me print this one out and ponder it for a while. But thankyou for the response. I think I can answer some of the questions, but others are going to test my ingenuity in dialectic. Thank you once again. I'll post within a few days (I hope - when up to ass in alligators ... but the damn things are only under control in imaginary space-time, not the real variety).

John G. [921125 22:39 CST gabriel@eid.anl.gov]

Date: Wed Nov 25, 1992 8:59 pm PST
Subject: RE: More testing

Who is the nut who keeps running all of those "tests?" Hi, Rick. Does the new address work as well on receive as on send?

Tom Bourbon

Date: Wed Nov 25, 1992 10:32 pm PST
Subject: Reply to Bill Powers of 921125

[Gabriel to Powers 921125, date and time 921125; 23:28 CST]

Bill, I'll reply in detail when I have a moment more than I do now. But a thought crossed my mind that I want to set down before I go to sleep. It moves the debate a little bit away from the classical PCT arena of discussion, but perhaps that will help understanding so please excuse.

There is a spectrum of human activity that we are each aware of in ourselves from circulation of the blood, or breathing which are life support systems that must go on whether we think about them or not, to inventing PCT or writing the Goldberg Variations.

There is not a well defined spectral line where on one side we can say we make conscious decisions and on the other side we don't, but suppose we abstract things and put one somewhere near the right place and agree not to argue further about where it is, at least for the moment.

On the conscious decision side, a lot of the time, there are discrete events we can call decisions, even if they are some time a making - perhaps almost a lifetime in some cases, and the spot where the decision is actually made is open to argument. Once again let's agree not to argue for the moment and proceed.

C. E. Shannon's 1949 papers put forth a rock solid proof, that given a channel carrying a known set of possible symbols (e.g. Morse Code), with a known frequency of occurrence for each (this can be made more general and more confusing, but let's not) and with each symbol having an independent probability of being chosen by the sender, then there is a

very well defined upper bound on the number of decisions a receiver can make per second on the basis of the received symbol stream.

Any communication - along nerves, by effectors, whatever, is subject to this bound, and it is called the channel capacity.

A channel carrying a truly continuous analogue signal with no noise has an infinite channel capacity in terms of discrete decisions per second. These kinds of channels have only relative capacities among each other.

No truly physical channel has infinite capacity. Even if you work at absolute zero so there is no thermal noise, there is still quantum noise from Planck's constant. So for a philosopher, there are no analogue channels, only discrete ones, and all channels have finite measurable capacity. And this is true for real engineers and real physicists too. But there are many cases, classical PCT is one, where the signal to noise ratio is so large that none of these information theoretic ideas carry any practical weight at all.

Just to make sure we have a shared percept, let me give a concrete example of a case where the information theory is important, that may make sense to all.

We all know the familiar 4KHz voice grade telephone line. How is it that carefully designed modems can transmit 56 Kbit/sec over a voice grade line, and why do BELLTEL need 64Kbit data rates for each voice grade line inside #5ESS?

Here is the argument. A good voice grade line has about a 256:1 SNR, so that only 256 levels of analogue signal can be distinguished. If there were more dynamic range there would be more levels, and if there were less noise there would be more levels. Thus, if we digitise the voice grade line to put it into #5ESS we only need an 8 bit A/D. Any more is wasted. For a 4 KHz line we need 8K samples/sec by the Nyquist sampling theorem, or we get aliasing. Hence the 64K bit/sec data rate in #5 ESS.

What about that 56 Kbit/sec modem. Well, that's a different story, even though it's constrained by Shannon, and 56 Kbit/sec looks suspiciously like 64 Kbit/sec with some allowance for slop.

It turns out that the 256 levels are what BELLTEL needs to be able to ensure that no perceptible distortion arises from the A/D, D/A conversions that go on at each end. The actual SNR of most voice grade lines allows only about 14.4 Kbit/sec, and that only with very elaborate means to defeat errors - a 14.4 Kbit modem has rather more computer in it than a MAC Classic. But now the 56 Kbit data rate is a puzzle. Was Shannon wrong??? It turns out that the answer is NO. You don't get the advertised 56 Kbit/sec with a totally random bit stream, you only get 14.4. But if you look at text or facsimile, there is a lot of correlation between one symbol and the next (what I was careful to rule out by making things simple early on) and you can throw away 3/4 of the data and still reconstruct things at the far end from what you know about the statistics. Of course in practice it isn't so simple - it takes another MAC classic, (not the hamburger) to encode things so as to safely throw away 75% of the incoming text and reconstruct it later.

So much for the digression. Revenir a Nos Moutons.

The whole purpose of the above was to support a claim that although the signal to noise and channel capacity issues don't matter for life support systems - these are much too important to risk atrial fibrillation just to save muscle mass in a rather small but vital component, they DO matter for Bach Preludes and Fugues, and for language. And they matter for politics and war and good government, all of which are my concern. In

particular they are at the centre of Martin's layered protocols, and my Kanerva spaces. It's important to understand them.

Well, I'm totally stonkered. It's after midnight and I started my day at 8AM. May I continue in a next installment?.

Let me say in closing that the fact that Martin, and I and probably Penni have something a little bit different in our minds from classical PCT does not diminish one whit our respect and affection for the guy who started it all. I don't think any one of us would have had the insight to see what you have seen, and the perseverance to see it through for more than 40 years. We all respect and revere the Einstein of the discipline, and we hope you will accompany us on the continued Odyssey.

I will hang up my hat sometime fairly soon, as will Bill C. But we have to cherish Martin who will be with the game for a long time to come, and Penni for longer still, and all the others whom I do not acknowledge so explicitly, or for many of them, at all. But a tip of the hat to Dag Forsell for his persistence, Greg, to Gary Cziko for gluing it all together, to Bruce and Rick for keeping us all honest, and all the others who are with us on the journey. And Particularly to Mary. My wife Marian says it's sometimes a bit trying to be married to somebody who has bright ideas at 4AM and has to talk about them RIGHT NOW.

John Gabriel (gabriel@eid.anl.gov)

Date: Thu Nov 26, 1992 2:52 pm PST
Subject: Complex percepts; decision etc.

[From Bill Powers (921126.1000)]

A Happy US Thanksgiving back. Too bad we didn't continue to get along with the Indians a bit better. PBS had a thing on Custer's Last Stand last night, enough to make one think of emigrating.

Martin Taylor (921125.1400) --

>As far as I can see, your model does use the same kind of computing
>elements at every level. All that changes is the perceptual input
>function, isn't it?

That's a pretty big "all." The input functions give us a world to live in, thoughts to think about, principles to live by, and so on.

>A configuration-level function has no delays, whereas a sequence-
>level PIF is basically a shift register, and a category-level PIF
>involves a catastrophe. If that's what you mean by non-identical
>computing elements, then I have no problem. To me, they are
>identical, in that they can all be described as multiple inputs to
>a PIF that produces a scalar value that is compared with a
>reference scalar value and the error transformed to a signal that
>goes down to lower ECSs.

A horse and a fish are identical in that they both eat and excrete. Between input and output, however, there are some differences that make a difference. I don't know how

these higher input functions work. I don't even know how configuration perception works. I don't think that understanding is going to get easier as we go to higher levels.

>I don't see how your version of reorganization can be guaranteed to be the
>only one.

I don't either. But I would like to get at least one kind of reorganization actually working in a model. So far the only kind we have actually got to work in a control-system model is YOUR kind (Tom Bourbon actually did it). It used the error signal as the basis for reorganizing the integration factor. Well, that's sort of my kind, too, as I think error signals could be classed as intrinsic variables.

>It isn't so much that the physics of the world and the nervous system
>determine what we can perceive (though this is true), but that they
>determine what percepts we can control, and therefore what percepts can
>be useful to the organism.

My point is that essentially ANY arbitrarily constructable percept is controllable. It doesn't have to be a "natural" form. You can control the sum of your distance from Durango in kilometers, 100 times the time of day, and 10 times the Kelvin temperature in Chicago. If you have a car. That wouldn't be a very useful controlled variable, but there's nothing in nature to say that this variable can't be controlled. I suppose that with practice you could even learn to perceive it. It's something other than the nature of the variable that determines whether or not we will learn to control it. And the environment doesn't constrain us to perceive it in any particular way. At least within limits.

The environment tells us what we have to do IF we want to control a particular variable at a particular level. It doesn't tell us (directly) what to want to control.

>... even within the same set of dimensions (the same set of lower-
>level perceptual signals, for example), there are an infinite
>number of possible sets of orthogonal representations at higher
>levels, especially when (as you point out) there are many different
>kinds of functions used in these representations.

We agree on that. It's interesting that as you construct more and more controllable variables out of a given set of perceptions, the possible controllable perceptions remaining get fewer and fewer because the set has to remain orthogonal. I wonder if this has something to do with the difficulty people have in changing system concepts that they're built up over many years. Maybe there's a point of no return, where you can't change anything without scrapping the whole orthogonal set and starting over. And maybe doing that could be fatal.

John (Stonker) Gabriel (921125.2328) --

>On the conscious decision side, a lot of the time, there are
>discrete events we can call decisions, even if they are some time a
>making - perhaps almost a lifetime in some cases, and the spot
>where the decision is actually made is open to argument. Once again
>let's agree not to argue for the moment and proceed.

It takes two to agree. I think that "decisions" most often are symptoms of conflict. Making a decision means not just noticing an error and acting to correct it, but imagining a second way of acting to correct it and then having to find a basis for

choosing which way to use. If you don't think up a second way, there's no decision to make.

When you walk across a room, you avoid obstacles without making any decisions. If the obstacle is to the right you move left, and if to the left you move right. But if you move left, you don't first consider moving right. This is a simple analog process that's built into my "Crowd" program; there's no considering of alternatives and choosing between them, although you could certainly put together a verbal description of the action that would make it seem as though dozens of decisions are being made. It's very interesting, in fact, when you know how the system is really organized, to see how many plausible verbal descriptions people can come up with. An "individual" can move into a pocket of obstacles, circle around, and finally double back and find a different way around. People will say "Oh, he sees that there's no way through here, so he's decided to back out and look for a place with fewer obstacles." Of course nothing of the sort is going on in the model.

When you walk across the room, the only times you have to make a decision is when avoiding one problem causes another one (if you go left you'll run into Aunt May, and if you go right you'll fall out the window), or when you're balanced on a knife-edge (the lamp is exactly directly in front of you). While we commonly attribute decision-making to people, I think that actual decision-making is relatively rare.

Mostly we just correct the apparent error; we don't consider doing something else first, and then choose which to do. Of course some people are chronic worriers and NEVER take an action without first considering five other things they could do instead. But that's just a sign of inner conflict.

>C. E. Shannon's 1949 papers put forth a rock solid proof, that given a
>channel carrying a known set of possible symbols (e.g. Morse Code),
>with a known frequency of occurrence for each (this can be made more
>general and more confusing, but let's not) and with each symbol having
>an independent probability of being chosen by the sender, then there is
>a very well defined upper bound on the number of decisions a receiver
>can make per second on the basis of the received symbol stream.

"Rock solid" as in mathematics, not as in modeling. The problems here are not in the mathematics, but in justifying the assumptions when you try to apply the mathematics to a real system. You can't justify the assumptions just by saying that they're necessary to the conclusions. You have to justify them INDEPENDENTLY.

So what is the "known set of possible symbols" when you try to apply this approach to human perception? A "symbol" can be any perception at all. Furthermore, most perceptions are continuously variable (even, near the boundaries, category perceptions). In the old days of information theory, everybody seemed to take it for granted that the basic unit of information was the neural impulse; either it occurred or it didn't, a nice neat binary code. But that assumption continued to be taken for granted long after the role of physical time was introduced to real neural models, after it became apparent that the time interval between impulses was important and that information was relayed in part through continuously-variable chemical concentrations. None of those considerations entered into the foundations of information theory as it was applied to the nervous system, at least the stuff I saw.

Then there's the "known frequency of occurrence" assumption. This is even weaker than the first assumption. You may be able to establish a known frequency of occurrence for texts emitted by many people over many circumstances, but you can't establish it for the next

thing a particular person is about to say. In order to do this rigorously, you would have to observe not only the symbol frequencies, but the environmental circumstances under which these symbols are used with a particular frequency, and the goals for which a particular person uses them, and produce a different frequency table for each possible circumstance and goal set. You'd have to see what is going on around a person, and know all the possible things that the person might want to say under those circumstances, and know what frequencies apply under just those circumstances. And you'd still be applying mass statistics to an individual event. So I don't see any practical way to apply the frequency-of-occurrence assumption.

Finally, what about "each symbol having an independent probability of being chosen by the sender?" Now we get the psychology of the individual sender into the act; if the sender is a painter, the probability of occurrence of numbers in conversation will be lower than if the sender is an engineer. If the sender is uneducated the probabilities will change; if the sender is given to cliches, likewise. There's nothing general that can be said about these probabilities that would be true of any individual.

These assumptions are mighty big if's. They define the playing field so narrowly that it's unlikely that the conclusions would apply to any real person. If the point is just to play out the game, fine. I'm all for entertainment. But if the point is to draw real conclusions about real people, the assumptions, it seems to me, are invalid on the face of it, so the conclusions don't apply. I think that in general mathematical analysts are far too hasty in adopting assumptions; they seem to adopt them so as to make the mathematics tractable, not so as to fit the actual world. Not you, of course, Stonker.

>.. there are many cases, classical PCT is one, where
>the signal to noise ratio is so large that none of these
>information theoretic ideas carry any practical weight at all.

I'm grateful to you for saying that. Maybe my preceding remarks are redundant.

Your nice signoff has me worried. Are you going away? Please don't.

Best to all, Bill P.

Date: Thu Nov 26, 1992 2:52 pm PST
Subject: Continued Discussion w/Bill P.

[To Powers, From Gabriel 921126 12:21 CST]

Dear Bill

This is a very quick post, just making some statements but not supporting them because I have other work which must be done today. Thus it gives an idea about where I am coming from and where we agree (lots of places) and where we disagree (some things I think we can resolve) but without chapter and verse.

We agree on:-

Feedback.
The ECS as an abstract concept.
The PCT hierarchy.
That actions spring only from perceptions.

I think that on reflection we will agree

We can talk about a chain:- perception -> decision to act -> action, at least for conscious activity. Perhaps we can agree about conscious activity in mammals, and about some kind of analogue of conscious activity in "The Moths and the Bats", and in stick insects.

At present I think we disagree about:-

The topology of the feedback network at high levels in the hierarchy.

The importance of Information Theory.

A "Universal ECS" rather like an OP AMP. I think that the abstract ECS is different at different levels of the hierarchy, once you need to talk about input/output path topology. But if the paths are trunks, not single signals then perhaps we don't disagree much.

The possibility of talking about state transitions as models of decision, although perhaps if I say that a decision is a minor reorganisation to shift the focus of attention, we may reach agreement.

"Universal Perceptions". I think that the perceptions in picking up a glass of water are different from the perceptions of a radar operator at a CRT, and are different again from those of an elector in a ballot booth. But I think the elector is not very different from a president deciding to declare war.

I think that agreement on these things might lead to profound and far reaching discoveries about the use of PCT in management, government and diplomacy. It also might let us scratch Dag Forssell's itch.

Sincerely and With Best Regards

John Gabriel [gabriel@eid.anl.gov]

Date: Thu Nov 26, 1992 2:52 pm PST
Subject: Re: Continued Discussion w/Bill P.

[Martin Taylor 921126 14:00] (John Gabriel 921126 12:21)

John,

I should be very surprised if Bill P agrees with your list of agreements. Specifically:

>We agree on:-

>

>Feedback. (yes)

>The ECS as an abstract concept. (a pretty concrete one)

>The PCT hierarchy. (depends on what you think it is)

>That actions spring only from perceptions. (NO NO NO)

>

>I think that on reflection we will agree

>

>We can talk about a chain:- perception -> decision to act -> action,

>at least for conscious activity. (NO NO NO)

The brackets are my guess as to Bill P's reaction. Anyway, they reflect my reaction. The NO NO NO comments are diametrically opposed to what I see as the heart of PCT, which is that behaviour is the control of perception. Perception does not lead to behaviour at all. It is the deviation of perceptual signals from their reference values that leads to action, and this action may be in opposite directions on successive appearances of the same perception, depending on the current level of the reference signal.

I don't think PCT has anything specific to say about conscious versus unconscious activity.

Martin

Date: Thu Nov 26, 1992 2:57 pm PST
Subject: Language Foul

[From Hank Folson (921126)]

Bill Powers (921122.2000) said:

>We have to consider what people use language
>for, not just try to explain language by examining words.

(penni sibun 921122.2300) responded:
>lots of us, both ``interactionists'' and more traditional
>psychologists, have been doing this for a long time.

Penni, your response is a good example of what my post (921122) was all about. Your response is true only to you and others who use a psychology based on other than control systems. To Bill and me and others who have a psychology based on control systems, it is wrong. In your world, the behavior of a person is interpreted as what the person is trying to do. In our PCT world, what the person is trying to do is control a variable, and the behavior that both PCT and non-PCT observers see bears little relationship to the controlled variable of PCT.

I perceive determined ignorance of PCT on your part. I don't think you would say that a Christian has to abandon her belief in Christianity to understand basically what a Muslim believes. You act as though you think an interactionist or more traditional psychologist has to abandon her beliefs to understand the basics of PCT.

If you understand (not believe, just understand) PCT, you never would have made that response. You would have understood that Bill & I were talking about 'controlled variables' in our PCT world (which you understand, but do not believe), and not the 'behavior' of your world (which you understand and believe in).

>language is fundamentally communication. wouldn't
>it be a good idea to start studying a language at its foundations?

To believers of PCT, the foundation is a control system controlling. To you, the foundation is something else. Is there common ground for working together?

Hank Folson

Date: Fri Nov 27, 1992 11:12 am PST
Subject: Re: Why Martin said NO NO NO

[Martin Taylor 921127 12:15] (Bill Powers 921126 20:40)

Bill's response to John Gabriel reminds me that there is an area I'd forgotten on which Bill and I agreed to disagree, and it is still there:

> One can understand and model behavior without getting into questions of
>information flow or quantity of information.

This disagreement is quantitative, not qualitative, because we agree that

> It [Information Theory] may have applications at the
>fringes -- signals near the lower threshold of detection, ambiguous
>situations in which the nature of a perception is not clear.

The disagreement is about:

> But in normal situations no such problems exist.

I believe that they do, as soon as you get any interesting distance up the hierarchy. It is easy to perceive whether your finger is importantly above or below a tracking target, not so easy to perceive, say, whether the sequence (program?) of docking the Queen Mary in New York without tugs is proceeding properly. (A distant cousin once did that). The perceptual input function involves all sorts of variables that interrelate, like tide state, river flow, ship mass and power, ship location and speed, wind, and probably lots of others. Each of those is perhaps measurable with some precision, but the relevant percept that is under control at a high level is much less quickly determined. And it is the rate of determination that is measured by information theory.

> Information theory is not necessary for designing or understanding the
>operation of a negative feedback control system.

Maybe not, but an understanding of information theory can sure help--a lot.

It's an old argument from last year, worth pursuing at some time, in connection with "slowing factors". Probably now is not the time, with other interesting threads going on here. I'd rather wait until we run into a lull on other matters.

I just wanted to say that the tattered old flag is still raised.

Martin

Date: Fri Nov 27, 1992 12:51 pm PST
Subject: Re: single vs multiple control: no problem to distinguish

[From Oded Maler 921127] (Bill Powers (921125.2000))

* >It's a relief. Otherwise if we accept "signal for every percept"
* >the only way to learn a new thing is to forget another...

* What, you mean that some people can learn something new without
* forgetting two other things?

Could you please remind me what we were talking about and who you are?

:-) --Oded

Date: Fri Nov 27, 1992 1:21 pm PST
Subject: Re: Complex percepts; decision etc.

[From Gabriel 921127 15:03 CST to Powers date/time N/A, rcvd 921127]

No, I'm not going away. All I was saying was that I am closer to retirement age than Penni, and closer than Martin, but not quite as much closer as I thought.

I've just received Martin's papers, and need to think. I DO tend to rush in before properly understanding the rules of the game.

But I'm certainly not going to stop playing. This is far too much fun. And thanx for your patience with a brash amateur. I will be interrupted from time to time by what Bill C. calls "the need to feed uncle", but don't get the impression that I've gone away. No way.

Your post gives me much food for thought. I'm still inclined to say that Information Theory is a good abstraction of the reality - at least as good as Newton was before Einstein, and that it's worth exploring its' quantitative and qualitative aspects.

I think you are absolutely right about all the possibilities in human communication that a particular recipient does not understand at any given time. I like Bach, but I don't understand Schoenberg. And if you have a continuous (i.e. non discrete) channel all bets are off and you have no absolute standards, only one channel vs. another with all the discrete channels having capacity zero.

Martin's Layered Protocols are important in establishing reliable communication. If you don't have his papers, they might help bridge some gaps. He's saying NONONO to some of my comments you have just replied to, but in different ways. I have not expressed myself well. But that's what the layered protoclcs should repair.

And I think your examples are better modelled by a continuous channel than a discrete one. There really are an infinite number of ways to circumvent Aunt May for all practical purposes.

If I'm a voter faced with a ballot, yes, there is a psuedo continuos channel in the write in vote. But unless I want to send the very special message "NONE OF THE ABOVE" (and so do a lot of other people) using it carries little weight.

So, perhaps both points of view apply each in a different context.

I'm really reluctant to give up Information Theory. It's almost always misinterpreted, and very easily supports paradoxes. I've seen lots of ecologists write about it, and all of them have talked nonsense. But your comments make sense. They seem to me to need proper answers, which will need a lot of careful thought. I have not read the linguistic psychological, or neurophysiological literature though and it will take time to do so. I hope to avoid the necessity because there is so much else to read and Uncle Sam does not support my interest in PCT (yet).

The difference between continuous and discrete channels is fundamental, and I think there may be a lot of meat in the 56Kbit/sec on a voice grade line example. This is particularly true because the actual modulation used on the voice grade line is phase modulation at least in some systems.

Best Regards John

Date: Fri Nov 27, 1992 1:25 pm PST
Subject: Re: Why Martin said NO NO NO

[To Powers from Gabriel 921127.1514] Bill Powers (921126.2040)

John Gabriel (921126.1221) --

Martin Taylor has correctly anticipated my responses to your proposed areas of agreement.

It's worth restating some basic facts about the "ECS."

Aha. We have agreement about the first of Martin's Layered Protocols, and I agree that whatever I may understand about PCT, we don't have a meeting of the minds, which is Layer 2 (if I understand Martin). Now we must recurse/change & iterate on the statements until we both believe the other understands what is being said. This is GREAT. I'm leaving transmit mode for a while to ruminate on what you have both said. Back in a while. Over and Out.

John

Date: Fri Nov 27, 1992 1:57 pm PST
Subject: Re: Why Martin said NO NO NO

[Martin Taylor 921127 16:30] (John Gabriel 921127.1514)

I think you are mixing up my layers with Pask's levels of understanding.

>Aha. We have agreement about the first of Martin's Layered Protocols, and
>I agree that whatever I may understand about PCT, we don't have a meeting
>of the minds, which is Layer 2 (if I understand Martin).

Pask's three levels may well be mapped onto three (possible successive) protocol layers, but which three will not be fixed. In principle, there is no lowest layer in LP, because no matter how physical you think you are, you can always break it down further. I usually, however, think of the acoustic wave or the visual pattern of light, shade, and colour, as being the lowest layer. Words, which are well below the level at which you received a correct message from Bill, are a loooong way above the lowest protocol layer, though for many purposes they might be the lowest layer of interest (e.g. in a speech recognition interface).

To show that words aren't critical to this argument, I can mention a visual example. I am a cricketer, and one of the key symbols to a cricketer is the weathervane on top of the Lords cricket ground in London. It is a figure of Old Father Time acting as an umpire to end the game. One of my friends had made a copy, and I wanted one. So I took the drawing from my friend, and gave it to a maker of weathervanes to copy. He did, but

only approximately, and his errors made it very hard to see the figure of Old Father Time with his scythe. The errors were numerically small, so he got Pask's first-level message. But he didn't get the second level because the scythe no longer looked as if its handle could be one piece of wood, among other flaws, and he most certainly didn't get the third level, which would have been the reason for this particular shape.

Martin

Date: Fri Nov 27, 1992 1:58 pm PST
Subject: Stable PIFs

[Martin Taylor 921127 16:50] (Bill Powers 921126.1000)

>Martin Taylor (921125.1400) --

>

>>It isn't so much that the physics of the world and the nervous system
>>determine what we can perceive (though this is true), but that they
>>determine what percepts we can control, and therefore what percepts can
>>be useful to the organism.

>My point is that essentially ANY arbitrarily constructable percept is
>controllable. It doesn't have to be a "natural" form. You can control
>the sum of your distance from Durango in kilometers, 100 times the time
>of day, and 10 times the Kelvin temperature in Chicago. If you have a
>car. That wouldn't be a very useful controlled variable, but there's
>nothing in nature to say that this variable can't be controlled.

But would it be useful when controlled?

Yes, any PIF of that kind could be made. I think I can describe (parts of) one that probably exists in the heads of millions of people: to perceive oneself with the forehead repeatedly touching the ground, knees and tops of the feet also touching the ground, shoeless, and with the toes pointing away from Mecca. The PIF of which this sequence forms a part can be very useful to control. Failure to control for it could lead to severe consequences for the states of one's intrinsic variables, if one were in some social communities.

I think that CEVs that we might think to be unnatural are likely to be ones that affect the values of our intrinsic variables only through social interaction. This (tends to) tie them into language, or at least into prescribed ritual. And it is probably specifically because such CEVs do not (naturally) affect the states of intrinsic variables that they can become elements of socially demanded ritual. When we control for them, we do not conflict with control of other "naturally" useful percepts.

So, we come back to the question of stable PIFs. I believe that under either your or my approach to reorganization, error in intrinsic variables leads to changes in PIFs and/or output linkages that are essentially random, and that the more the error, the more likely the change. If I happen to be controlling for "the sum of your distance from Durango in kilometers, 100 times the time of day, and 10 times the Kelvin temperature in Chicago," and I happen to have some error in intrinsic variables, that PIF is as likely as any other to be changed. Since it has (by assumption that it has no social implication) no effect on intrinsic variables when it is controlled to different reference levels, it has a good chance of being changed, along with other PIFs, until the intrinsic variables tend to stay near their references. If controlling a percept of this form does matter to the

intrinsic variables at what level it is controlled, then altering that PIF is likely to make matters worse, and there will be more reorganization.

In the end, quite a few irrelevant PIFs may survive, but most PIFs will exist because they produce percepts that are both (usually) controllable and useful. In my previous posting(s) I concentrated on "controllable" because without that, a PIF cannot be useful. But a PIF survives only by chance if it is not useful, and it wastes resources. Things that waste resources are not treated kindly by evolution, and I would speculate that there is a tendency to eliminate such PIFs, if only by the random-walk kind of process that is reorganization. Some remain (we call the corresponding actions "superstitious"), but most will make way for PIFs that are both stably controllable, and useful in keeping intrinsic variables near to their reference levels.

(But see the discussion earlier in the year about "The Bomb in the Hierarchy").

Martin

Date: Fri Nov 27, 1992 2:12 pm PST
Subject: Decisions Decisions.

[From Gabriel to Powers & Net 921127 15:57 CST]

Perhaps Decisions happen when you have an interest in the outcome, but have no hope of total control. E.g. as in war and in politics.

But in picking up a glass of water I agree you don't really decide to pick it up, except if you want to use decision theory to choose between beer and water. I.e. decisions have to do with long term payoffs - Do I want to drink beer and be hung over in the next round of discussion on CSGNET?.

Long term payoffs are always uncertain, and affected by the unpredictable actions of others, which you only estimate by assuming that the time series from the past is stationary. If you discover (painfully) that it isn't then you change your doctrine or political strategy. The stationary time series is just another name for the Shannon ensemble.

So. Perhaps another aspect of decisions is that they involve some reasonably small number of discrete choices, and presumably reorganisation (I don't have a reliable percept for reorganisation, but perhaps a vague one will do).

Also Information Theory is at the heart of Error Correction, which PCT may take for granted by saying "assume there is an error signal."

DAMN I meant to stay in ruminative mode, but this was a rumination too good to miss sharing it. BELCH!!!

John (gabriel@caesar.eid.anl.gov)

Date: Fri Nov 27, 1992 2:15 pm PST
Subject: Re: Why Martin said NO NO NO

Curiouser and curiouser. I don't have any response yet, but I'm drinking it all in. Thankyou. Msg received.

John

Date: Fri Nov 27, 1992 2:38 pm PST
Subject: Symp: New directions in control theory and applications

[Martin Taylor 921127 17:30]
Here is a cross-posting from comp.ai.neural-nets.

**
** IEEE Meditteranean Symposium on new Directions **
** in Control Theory and Applications **
**
** Chania, Crete Greece **
** June 21-23, 1993 **
**
** Organized by: **
** IEEE Control Systems Society Greek Chapter **
** The Laboratory of Automation, Technical University of Crete **
** Participating Society: IEEE Control Systems Society **
**

SCOPE

=====

The aim of this IEEE Conference is to bring together CONTROL SYSTEMS RESEARCHERS AND PRACTITIONERS (systems people) working in Academia, Government and Industry so they can discuss the latest developments in modern control. Emphasis will be placed on INTELLIGENT and MODERN CONTROL methods and their applications

TOPICS

=====

(I) METHODS

Modeling	Linear and Nonlinear Systems
Intelligent Control	Neural Networks
Discrete-event systems	Adaptive Control
Stochastic Control	Numerical Methods
Fuzzy Logic	

(II) APPLICATIONS

Robotics	CIM
Quality Control	Computer Control
Real-Time Software	Man-Machine Systems

Date: Fri Nov 27, 1992 2:47 pm PST
Subject: Reorganization; Information theory

[From Bill Powers (921127.1500)] Martin Taylor (921127.1650) --

>In the end, quite a few irrelevant PIFs may survive, but most PIFs
>will exist because they produce percepts that are both (usually)
>controllable and useful. In my previous posting(s) I concentrated
>on "controllable" because without that, a PIF cannot be useful.
>But a PIF survives only by chance if it is not useful, and it
>wastes resources. Things that waste resources are not treated
>kindly by evolution, and I would speculate that there is a tendency
>to eliminate such PIFs, if only by the random-walk kind of process
>that is reorganization. Some remain (we call the corresponding
>actions "superstitious"), but most will make way for PIFs that are
>both stably controllable, and useful in keeping intrinsic variables
>near to their reference levels.

In the above quote, I first deleted the hard carriage returns that came with it, then inserted a > before each line, then ended each line with a hard carriage return. My left margin is 0 and my right margin is 72. Did it come through with the >'s displaced? If the receiving party has the right margin set less than 72, the lines will end prematurely, and the next line will be a short continuation with a hard carriage return; but the long lines will have a > in the leftmost character position. If the >'s are showing up in the middle of lines, then something or someone dastardly is removing my hard carriage returns between the screen I'm looking at and the one you're looking at.

As to the above quote, we can nail that down. We have the same picture. I agree. You agree that I agree. I agr...

What remains to be resolved, and it won't really be resolved without some heavy modeling, is (a) how reorganization affects the reorganized system, (b) what conditions initiate and terminate reorganization, and (c) what constraints determine possible outcomes of reorganization. I am satisfied to leave those matters for future development. We can try out lots of ideas and see how they work.

John Gabriel (921127.1503) --

>I'm still inclined to say that Information Theory is a good
>abstraction of the reality - at least as good as Newton was before
>Einstein, and that it's worth exploring its' quantitative and
>qualitative aspects.

My only real objection to information theory arises when it's used to explain things that it describes but doesn't explain. There's a difference between noting that in some system design the principles of information theory are borne out, and claiming that the system design works BECAUSE OF information theory. This is a lot like my objection to Ashby's "Law" of Requisite Variety. After a control system has been designed to operate satisfactorily, you can note that the "variety" of the system's output matches the "variety" of disturbances of the controlled variable. But the control system does not operate properly because variety has been matched; variety matches because the control system operates properly.

The detailed design of a system does not have to employ information theory; generally the constraints on a system design are physical, set by the limits of components and materials. Telephone engineers do not deliberately limit the bandwidth of ordinary telephone lines to 4KHz. That's just the most you can get without going to a lot of expense. You can get more if you want to pay more. Bandwidth isn't limited by information theory, but mainly by line capacitance and inductance and transformer core losses as a function of frequency. Information theory can be useful in telling you what you can

expect to get through a system designed in a particular way, but it won't help you discover how to increase the inherent bandwidth. That's an engineering problem.

The explanation for the bandwidth limitation is in the physics of long- line transmission. A description of the consequences of that physical limitation can be found in information theory. But information theory does not explain why the components have those limitations. Capacitance and self-inductance are not aspects of information theory.

Best to all, Bill P.

Date: Fri Nov 27, 1992 3:04 pm PST
Subject: Decisions

[From Bill Powers (921127.1600)] John Gabriel (921127.1557) --

>Perhaps Decisions happen when you have an interest in the outcome,
>but have no hope of total control. E.g. as in war and in politics.

A hierarchy of control exists, I think, mainly to avoid having to make decisions. When you're going out the door to go the drug store and pick up a newspaper and your wife calls out "Don't forget you have to pick up Dolly at the train station," you would have to decide between getting the newspaper and picking up Dolly if you weren't a hierarchical control system. But being such a system, you can use a common set of lower-order behaviors to accomplish more than one goal at a time. You drive to the train station, and while you're waiting for Dolly to arrive you buy a newspaper THERE. Brilliant. If you couldn't alter lower-level goals to achieve multiple higher-level ones, you would have decide whether to give up the newspaper or meeting Dolly.

For a hierarchical control system, the usual answer to "Which goal should I achieve?" is "Both."

Best, Bill P.

Date: Fri Nov 27, 1992 4:11 pm PST
Subject: Re: Decisions

[From John Gabriel 921127.1707]
>[From Bill Powers (921127.1600)]

>John Gabriel (921127.1557) --

>>Perhaps Decisions happen when you have an interest in the outcome,
>>but have no hope of total control. E.g. as in war and in politics.

>A hierarchy of control exists, I think, mainly to avoid having to make
>decisions. When you're going out the door to go the drug store and pick
>up a newspaper and your wife calls out "Don't forget you have to pick up
>Dolly at the train station," you would have to decide between getting
>the newspaper and picking up Dolly if you weren't a hierarchical control
>system. But being such a system, you can use a common set of lower-order
>behaviors to accomplish more than one goal at a time. You drive to the
>train station, and while you're waiting for Dolly to arrive you buy a
>newspaper THERE. Brilliant. If you couldn't alter lower-level goals to

>achieve multiple higher-level ones, you would have decide whether to
>give up the newspaper or meeting Dolly.

>For a hierarchical control system, the usual answer to "Which goal
>should I achieve?" is "Both."

Yes, I agree, and I think you've captured a very important issue. You do a lot of what ifs (i.e. run models) to see how to maximise the weighted sum of the truth values of a set of goals. And if you find a strategy that achieves a whole lot of important goals that's great. Achieving a lot of unimportant ones and missing the big one is not so great.

Now, what if your models are unreliable? A drunk driver may run into you on the way to the station, the adversary may have mobile reserves you did not know about. All this gets into a discussion I have with Tom Baines about granularity, time scale, predictability, and signal to noise in estimated outcome.

I don't have all the ideas sorted out yet, but it seems to me that without reliable estimates of the effects of attempted actions on the environmental variables, perhaps we have to extend the model of PCT a bit. And possibly this is what Martin is thinking of too in his discussion of docking the Queen Mary.

Aha, this is an issue about different bandwidths in different feedback paths. If things can go to hell beyond retrieval before you know it, then the model of simple control needs extensions for those cases.

Tom Baines has a wonderful example, I think from van Creveld or from Keegan about the Greek general who was sitting down to a feast because his center had routed the enemy, when a messenger arrived to say that his right and left wings had been driven from the field and the enemy was at the gates.

Well, for a guy who was going into rumination mode, I'm doing an awful lot of talking. Perhaps these are the ruminations, and putting them up for other people to look at finds the mistakes more quickly than I would on my own. Thankyou all for your patience and helpful replies. The NET is a great place. Much better than the classical referees. Perhaps in due course nets and archives will replace journals. I think we might progress faster if they did, and the only stuff published on paper was things like Reports on Progress in Physics, or the ACM tutorial journal.

Sincerely John (gabriel@caesar.eid.anl.gov)

Date: Fri Nov 27, 1992 8:07 pm PST
Subject: Decisions

[From Bill Powers (921127.2030)] John Gabriel (921127.1707) --

>You do a lot of what ifs (i.e. run models) to see how to maximise
>the weighted sum of the truth values of a set of goals. And if you
>find a strategy that achieves a whole lot of important goals that's
>great. Achieving a lot of unimportant ones and missing the big one
>is not so great.

Actually I don't do very much of that, although I understand that a lot of people believe that they do. I work mostly in the reorganization mode -- if something doesn't work, try something else. Things just gradually fall into shape. Maybe that's my problem.

Docking the Queen Mary from the helm isn't really a good example of this, because it's an analog problem (Martin note). Actually, docking the Queen Mary ought to be easier than docking a small motorboat, for the same reason that balancing a broomstick on end is a lot easier than balancing a foot ruler on end. It isn't really necessary to take wind, tide, and other such external causes of disturbances into account, if you've put together a good control system that can directly affect the controlled variables themselves. A huge ship can't do much that's outside the bandwidth of an ordinary neural control system.

Back in the 50's people used to run problems like that on analog computers. At 60 iterations per second.

The basic PCT objection to planning (as it's ordinarily done) is that the output of the plan is a prescription for actions that are predicted to have certain effects. A control system doesn't assume that a given action will have a particular effect. Instead, it monitors the effect, and is prepared at all times to alter the action if the effect deviates from what is desired.

When the effect is in the future, the PCT approach would be to monitor progress toward the effect, and control that. In complex systems, once you have done the best you can to set up negative feedback control systems, you can think about adding predictions and so-called "feed-forward." These additions will probably give some improvement of performance, but only if there is an underlying control system to take care of the unexpected.

>I don't have all the ideas sorted out yet, but it seems to me that
>without reliable estimates of the effects of attempted actions on
>the environmental variables, perhaps we have to extend the model of
>PCT a bit.

I want to reemphasize that control systems do NOT depend on "reliable estimates of the effects of attempted actions on environmental variables." They do not attempt any estimates at all. This is one of the principal differences of the PCT approach from conventional ones. All that a control system has to do is to be able to sense the state of a controlled variable, and to affect the state of the controlled variable. The feedback process adjusts the output on the basis of deviations of the controlled variable from its intended state until the deviation is zero or sufficiently small. Then it keeps the deviation small. This adjustment occurs continuously; this is not a process of producing an action, waiting to see its effect, estimating a correction in the action, producing another action, and so on. That is the sort of model that has been produced by people who are wedded to the old paradigm.

It is always possible to start with the PCT core design and add features to cope with unusual situations -- for example, situations in which there are forewarnings of disturbances of sufficient accuracy to allow precalculating a move to oppose them. In most cases there are no such forewarnings. When there are forewarnings they are seldom accurate enough to allow constructing a quantitatively appropriate action to oppose them. Without quantitative accuracy, the blind opposition is just as likely to make the final error worse as better.

There are limits to what control systems can do. When the situation demands more than can be done with control, the result is going to be undesirable no matter what substitute is

used. So there is a very strong motivation to figure out how to design control systems to take the place of conventional planning.

Best, Bill P.

Date: Fri Nov 27, 1992 9:42 pm PST
Subject: Re: Decisions

[From Gabriel 921127 23:07 CST] Powers 921127

Yes, I understand what you are saying, as far as it matches my own experience as a designer and builder of complex control systems.

But I don't entirely agree. What I hear you say, is that at large times compared to the inverse system bandwidth, the error signal tends to zero. And I agree entirely, at least for linear systems.

But suppose you are the captain of a large oil tanker that cannot stop in less than a mile, that you have no charts for where you are, and whatever means you have to recognise rocks has 1/4 mile effective range. When you try to dock, as the EXXON VALDEZ?? did you have some difficulties. One solution is to slow down to 1/4 knot, so that you can stop by ordering full speed astern. But even there you must be careful not to overshoot if there are rocks behind you too, so you need a good model of the dynamical system.

And if you are a president running for re-election, a general fighting a campaign, or a president trying to run the executive branch, your opposition will move the rocks around into your path if they can. And if you are only running 1/4 knot, and there is an 8 knot tide, the tide is controlling instead of you.

This is the real problem that confronts a president, a CEO, or the commander of a campaign. You can see it most clearly trying to trade on Wall St.

But whoever you are, and whatever you do, you can only act on your perceptions. That is the central message of PCT, at least for me. Because that is true no matter what the system bandwidth may be, and so the staff to the president, the CEO, and the general all have the same job. To observe the environment in more detail than the Cdr. has time to do for him/her self, select the things that need attention NOW and tell the Cdr.

This is a problem in Information Fusion - new term, coined by Bill C. And it would not have come into existence without PCT. And he and I and the others who are concerned with good government or avoiding avoidable wars and winning unavoidable ones are absolutely indebted to your insight and persistence. And we thank you from the bottom of our hearts. Even if we disagree with you about the little things.

And I am still not confident that I understand everything you say, in fact I doubt I do, and I also doubt I understand Martin properly. But one day I will. That's a reachable limit point because we none of us make irrational or unreasonable changes, and there is an intention to reach agreement. Now that's a PCT feedback loop we can control if we work hard enough and long enough.

Sincerely and with Very Best Regards

John Gabriel (gabriel@caesar.eid.anl.gov)

Date: Sat Nov 28, 1992 12:44 pm PST
Subject: Decisions; hierarchical organizations

[From Bill Powers (921128)] John Gabriel (921127.2307) --

I understand and appreciate your intention to apply the principles of PCT in these difficult fields of behavior. Sometimes in the heat of grappling with ideas I forget to say how grateful I am to you, Martin, and all the others just for being here and joining your efforts with mine. Also, I know that I am often guilty of trying to make complex problems seem simpler than they are. I don't mind being reminded of that.

>But whoever you are, and whatever you do, you can only act on your
>perceptions. That is the central message of PCT, at least for me.
>Because that is true no matter what the system bandwidth may be,
>and so the staff to the president, the CEO, and the general all
>have the same job. To observe the environment in more detail than
>the Cdr. has time to do for him/her self, select the things that
>need attention NOW and tell the Cdr.

Yes, that is the central message. As long as that is kept intact, the rest will eventually sort itself out.

The concept of hierarchical control is also important, and is perhaps even more difficult to bring into the orbit of a specific problem like yours. When you deliberately set up an organization to emulate a control system (it can't actually be a control system, but people can agree to pretend, for a particular purpose, that it is), some of the principles of hierarchical control can be helpful.

In the human hierarchy, there is no duplication of function among levels for the simple reason (if my idea is right) that each level handles a different logical type of variable. This means, for example, that if relationship perceptions need controlling, they are controlled at the relationship level, and higher levels need know nothing about how to control relationships. The higher level only needs to be able to specify that a certain relationship be realized in perception, and the lower systems see to it that this happens. And generally the higher level does not require that a specific relationship be brought about -- only that some combination appear that satisfies the pattern under control at the higher level.

In an organization, the parallel problem is that of what perceptions and goals are appropriate to each level. It is not just that the subordinates tell the CEO what things need attention now, but that they provide information that is relevant to the kind of thing a CEO pays attention to. This should not be the same kind of information that a subordinate pays attention to; it should be a synthesis of what many subordinates are controlling for, a variable of a new type with which no one subordinate can be concerned.

The subordinates should never be told what to do. They should be told what to perceive. How they arrange their own worlds so as to perceive the specified variables in the specified states is up to them. It must be up to them, because only they experience the disturbances and other difficulties that pertain to the actual process of control at their own level, and only they are in a position to correct errors in the least possible time. If a higher system tries to direct the ACTIONS of a lower system, the result will be conflict, sluggish control, and even instability of the whole system. The CEO simply does not have, and can't have, the proper organization for controlling the way goods are purchased by the Purchasing Officer, or how the mechanic mounts a tool on a lathe.

Instead of thinking of the hierarchy in an organization as a hierarchy of command, it should be thought of as a hierarchy of specialists. Different people can specialize in controlling for different levels of variables, so they emulate the functions of a particular level of control in the organization. One trick in designing an organization, it seems to me, would be to define the levels of control so they are really of different logical types. Furthermore, within a single level, the trick is to design the various jobs of the same logical type so that the individuals control for variables that are orthogonal, meaning that each person's controlled variable can in principle be in any state without conflicting with any other person's controlled variable of the same type. When orthogonality exists, the higher levels have the most freedom to vary the goals they set for the lower levels. Also, orthogonality makes it easy for each subordinate to achieve the requested perception without running into conflict with others at the same level; then the higher levels can be assured that the goals they specify will actually be realized.

Most of the organizations I have worked in, starting with the United States Navy, seem to have made an attempt to work out principles like these, although without any systematic theory behind the attempts. All of them have contained serious sources of conflict, so that for one person to do the specified job, someone else has to be frustrated. All of them have confused functional position in the hierarchy with social position or worth. All of them have had to rely on coercion to greater or lesser degrees. And all of them have violated the hierarchical control principle that a higher level system should not try to direct the actions of lower-level systems.

Just through random reorganization, however, many of the principles of hierarchical control seem to have been worked out or at least approximated. Organizational protocol, for example, prevents perceptual and reference signals from skipping levels. There seems to be a growing realization that coercive command does not work very well. There is an apparent tendency toward giving subordinate levels more autonomy and more ability to vary actions as judged necessary at a given level. Delegation of authority is being tried more -- although I haven't seen anybody talking about logical types of controlled variables, which is what delegation should be about.

It would be fascinating to try to design an organization on these principles. As I often tell people thinking along these lines, however, it's important not to forget that a social emulation of a hierarchy of control is not like the hierarchy in an individual. In an individual, one level of control is concerned with one level of perception, and knows nothing of other levels. Each component in the social hierarchy, on the other hand, actually functions at all the possible levels from neuromotor coordinations to policy setting. Just listen to the debates over company policy that go on in the blue-collar lunchrooms, or watch the VPs playing handball. It's impossible to confine a person's perceptions and actions to a single level of function in a social hierarchy. The best that can be done is to obtain agreement to and understanding of the organizational function the person agrees to fulfill.

All that probably has nothing to do with your interests, but it's fun to think about.

Best, Bill P.

Date: Sat Nov 28, 1992 8:41 pm PST

[Martin Taylor 921128 23:30]

(Bill Powers 921128 1330 by the receipt header, to John Gabriel)

Nice discussion of organizational structure.

>The subordinates should never be told what to do. They should be told
>what to perceive. How they arrange their own worlds so as to perceive
>the specified variables in the specified states is up to them.

This is fine if the superiors believe the subordinates to be competent, but in contrast to the situation in a control hierarchy, the superior may know what ACTIONS the subordinate should take in order to get the desired perception, better than the novice subordinate knows. The boss often must be a trainer, providing what in a hierarchy of ECSs would be guided reorganization. We don't think that can happen in a control hierarchy, but in an organization, it can. I think that makes a difference. In a well running organization in which everyone knows their job, I'd agree with just about everything in your posting. But to get to that state is different from reorganization in a control hierarchy, because people are constantly moving jobs, and because the boss really does often know the junior's job, having been there. No ECS can be in that condition. The boss ECS has never been in the subordinate's place.

But once the subordinate knows the job moderately well, then of course the boss should just ask to be provided with certain percepts (I think).

> If a higher system tries to direct the ACTIONS of a lower
>system, the result will be conflict, sluggish control, and even
>instability of the whole system. The CEO simply does not have, and can't
>have, the proper organization for controlling the way goods are
>purchased by the Purchasing Officer, or how the mechanic mounts a tool
>on a lathe.

They say that one of the problems of the Vietnam war was that the generals could see and command at very fine levels of detail, directing the actions of the soldiers in the field. Not a great idea.

Martin

Date: Sat Nov 28, 1992 10:13 pm PST
Subject: Re: Decisions; hierarchical organizations

[From Ray Jackson 921128.2300]

Bill Powers (921128) on John Gabriel (921127):

Bill,

As you probably remember, my job is to develop an empowering culture (whatever that is), in a major electronic manufacturing plant here in the Phoenix area. The biggest challenge I have is to give meaning to the word EMPOWERMENT, which is probably the most overused and least understood business term in history. From that perspective, I find HPCT the missing link which gives meaning to the word and, as I've said before, actually explains the essence of the value to be found in empowerment.

>The subordinates should never be told what to do. They should be
>told what to perceive. How they arrange their own worlds so as to
>perceive the specified variables in the specified states is up to
>them. It must be up to them... If a higher system tries to direct

>the ACTIONS of a lower system, the result will be conflict, sluggish
>control, and even instability of the whole system.

>Instead of thinking of the hierarchy in an organization as a
>hierarchy of command, it should be thought of as a hierarchy of
>specialists.

You're absolutely right, further illustrating the critical role of (valid) communication and (valid) education in developing an Rorganizational variableS. Unfortunately, these two areas are notorious corporate weaknesses, even in good companies.

>One trick in designing an organization, it seems to me, would be to define
>the levels of control so they are really of different logical types.

There's a fascinating and innovative slant on Organizational Development, and one that makes a tremendous amount of sense (usually not a major consideration in the business world). OD is generally micro-minded, focusing on areas and links, rather than a broader systems approach. I guess this is due to that Industrial Age, S/R mindset.

By the way, before Thanksgiving, I was approached by Grand Canyon University to teach night classes for their MBA program next fall (Organizational Behavior and Organizational Development). I do speaking engagements with them every semester, and they're familiar with what I do; I can't wait to get my hands on the material to weave PCT through it. Gosh, with Tom Hancock full-time in the College of Ed, and myself as an adjunct in the College of Business, there's no telling what we might be able to facilitate there. (GCU a PCT stronghold?)

>There seems to be a growing realization that coercive command does
>not work very well...

Of course, that's the refreshing part of what I do. Often times, the growing realization is an error signal produced by a genuine concern for success, on both the organizational and personal level. This creates a very positive learning and growth situation. On the other hand, sometimes the error signal is generated to answer a mandate or follow a fad, and well you can imagine the outcome there.

>Delegation of authority is being tried more -- although I haven't
>seen anybody talking about logical types of controlled variables,
>which is what delegation should be about.

You probably haven't. Currently, effective delegation techniques revolve around clarity of expectations and conscientious follow-up. Also, keep in mind that delegation is usually resorted to out of necessity (doing more with less), and most individuals want to get rid of some of the work without giving away any of their power; this is especially true when people are scared (as most are today).

>It would be fascinating to try to design an organization on these principles...

You're right...but it's not as fascinating as it is necessary in today's business environment. Also, despite the obvious competitive benefits for an organization to pursue this type of culture, they do at the risk of enhancing the quality of life for their workforce; it's the true win-win situation, which is why this works. Potentially, as more companies embrace these principles, it can have far reaching implications for society as a whole.

>All that probably has nothing to do with your interests, but it's
>fun to think about.

Ah, but it had everything to do with my interests (and Dag's, and others I'm sure...).
This was the first time I've heard you speak on the Net on such topics. Thank you.

Best Regards, Ray

Date: Sun Nov 29, 1992 12:53 pm PST
Subject: Social system, Error-free

[From Dag Forssell (921129.12.15)]

Re: Decisions; hierarchical organizations, Ray Jackson 921128.2300

>This was the first time I've heard you speak on the Net on such topics.
>Thank you.

For a more extensive treatment, see "CT Psychology and Social Organizations" in LCS II,
page 91-127.

Re: Error-free programming.

Thanks to John Gabriel, Hans Blom and Martin Taylor for notes and suggestions. I put
together a report, and passed it on. Whether it falls in fertile soil I cannot say.

My (younger) daughter reprimanded me for trying to interest others in my passion, and
wants me to be more sensitive. She says that she perceives her friends to feel compelled
to be polite to me and suffer my chatter.

I find that I cannot respond to a polite inquiry from any guest as to what I do, without
my sensitive daughter thinking that I am abusing them, even when she only catches a
glimpse from across a crowded room. An indication of professional interest seems no
excuse for me to continue a conversation and offer information. Is there not an old
saying about being a prophet in your own country?

I feel blessed to have something to be passionate about.

We are talking about it and exchange views. I pointed out to her that this is only her
perception, (in large part based on incidents she partly misunderstood many years ago)
and asked her to think carefully about her perceptions, before she acts on them. As with
promoting my program, I find that I have to be extremely careful about how I express my
convictions and indicate the implications. (She let me know that to "teach executives how
people function" implies that everyone who does not already know is an idiot). PCT helps
resolve misunderstandings even about PCT. It was good having the kids home. What I am
relating here was the only - and quite minor - irritant. Now they are back to job/school
again.

Meanwhile, I am making good progress on the promotional front, collecting names and
preparing phone questionnaire to ask people who have had quality training what their

concerns are, and what clarifications they would like to make the quality "implementation" easier.

Best to all, Dag

Date: Sun Nov 29, 1992 1:02 pm PST
Subject: Re: Decisions; hierarchical organizations

Dear Bill

Thankyou very much for your post on Decision; hierarchical organisations. Far from being not what Bill and Martin and I are worrying about, I think your comments are right on the mark, and have new insights for us. I (and I am sure, the others), welcome you warmly. The idea of new logical variables at each level is new, at least to me, and is very close to some of the things I have been grappling with in mapping semantic networks onto what an organisation can be said to "understand". Thanks again. And your Navy experience is exactly matched to what we need in understanding military organisation. Eventually I'd like to look at some of the ways we can use PCT in large corporations, but the military is an interesting laboratory because the nominal rules of how to work are set forth in Field Manuals, and there are pretty reasonable transcripts of exercises where you can see what people actually did, if not what they thought. So you see the analogue of some of the things we cannot observe in transmission between levels of the hierarchy in people.

So, I agree with you warmly, and I think completely, as far as I understand my own problems and your statement of them. I'm looking forward to an enjoyable collaboration to the extent that you have time and patience.

Sincerely John

Date: Sun Nov 29, 1992 1:49 pm PST
Subject: standards in counseling

from Ed Ford (921129:1445) Rick Marken (921124.1000)

>but you are dealing with some enormously precient people;

I deal with all kinds, from those in my private practice (including minority and poor, who pay on a sliding scale) to very wealthy; from the very young to the very old, in various types of residential settings; from corrections to school. Maybe the common factor is that we all live in Arizona.....

>apparently, they are able to know what the state of the world
>(disturbances) will be when they set their standards, rules, and
>criteria so that these will be appropriate and allow them to reach
>their goals. How can they do this when the disturbances they will
>actually encounter are unpredictable and, often, undetectable? I
>thought that PCT made it clear that the only appropriate settings for
>ANY references are those that, when the outputs resulting from these
>reference settings are combined with prevailing disturbances, produce
>the intended perceptual results.

It is easy to say things in words in the theoretical realm. I wish you would use several examples, it would be much easier for me to understand and to deal with precisely what you are saying. In any event, I will try to respond to what you've said.

I think the answer is in the fact that any time we have a goal (reference), and we attempt to achieve this goal, the standards or criteria we set can be set for many reasons, many having to do with other references that interconnect or interrelate to the main reference I have. Whether I am trying to decide on which university to attend, or a young woman to marry, or to drive on a free way, or to exercise, or to eat "healthier" foods, or where I want to live, or just to call a friend, all these references are going to involve my making choices which are going to involve other important references. I might set some standards for the kind of woman I want to marry, but in my attempt to satisfy this goal, I might have to adjust my standards if my choices reject me. I might have certain standards for the way a happily married couple should live, obviously those of us who are (or were) married have found a constant adjusting of standards very necessary to meet the "happy and warm, loving relationship" goal. As a vegetarian, I have very strict standards for what I eat, but I don't try to impose those standards on those with whom I live. Often, when asked to dine at the home of a friend, I willingly adjust some but not all of my standards (I'll eat some cheese but never meat or fish). I have been successful on maintaining a no smoking policy in my house by asking visitors who must, to please smoke outside.

All the perceptual variables with which I am trying to deal can be controlled only by satisfying all the other interconnected references as well as the one I'm trying to satisfy. Standards can describe in specific terms of the kind of variable you are controlling for, they can also describe the outer limits you are willing to go to to reach or achieve your reference, including how much of the disturbance you are willing to tolerate. Standards can also be tied to another reference that is definitely interconnected or interrelated to the present reference which you are trying to satisfy.

>I said (921121) you teach people how to use their own system, to set
>their own goals, their own standards upon which they can make choices,
>because PCT teaches me that that is how the system is designed.

>then you said (921124) Now the question is how does one do this. I
>don't know how but I don't think it's by getting people to articulate
>"appropriate" settings for references; it must have something to do
>with figuring out what the person is conflicted about; conflict being
>the main reason for lack of control.

I agree, it certainly is not getting people to articulate "appropriate" settings for references. Rather, it is to have them articulate and then evaluate their present settings for their references and to see if these particular settings are the most efficient or best settings and are these the best standards for helping them reach their goal.

You are absolutely right when you said "it must have something to do with figuring out what the person is conflicted about" - but I would say it would be more accurate to state "it must be something to do with helping the person figure out what his/her conflict is all about." IT IS THE PERSON WHO HAS TO DISCOVER THE SPECIFICS OF HIS/HER CONFLICT AND THE ESSENTIAL ELEMENTS WITHIN THE CONFLICTING AREA THAT NEED TO BE EVALUATED, INCLUDING THE REFERENCES AND PRIORITIES THEY'VE SET, THE STANDARDS AND CRITERIA THEY'VE ESTABLISHED, AND THE CHOICES THEY'VE MADE AND WHETHER ANYTHING IN THIS CONFLICTED AREA NEEDS TO BE CHANGED OR ALTERED TO REDUCE THE CONFLICT. Don't you see, this is what I've been trying to do. It isn't the counselor that has to discover all this, it is the

person who is having the conflict that has to discover this. All the counselor is doing is helping or teaching the person to better use their system more efficiently and effectively.

>So it's not that I think that you are trying to impose YOUR values -
>it is that you are suggesting that there are values that are right for
>the client...

No I am not. I am suggesting the client find the best standards or criteria that work best for him/her in the situation in which they find themselves. My job is to teach them how to use their hierarchical system as suggested by PCT. When I ask them about their various levels, I am actually teaching them to think "level wise" and to think about the interconnectedness and the interrelationships involved. From that they are better able to articulate to themselves (and to me) the specifics of what is going on in their world. The more they understand how their living control system works, the more they are able to use it to their own advantage. There in lies the beauty of PCT and especially the levels. When the levels are understood in light of how we function, they become much more useful to us and our ability to manipulate our own system to our own advantage is enhanced so that we can satisfy our own internal goals and thus eliminate or reduce conflict to a point where we can live with it.

>it depends on what they are trying to achieve at a higher level (which
>I think you clearly understand) AND it depends on prevailing (and
>unpredictably changing) circumstances - so that the setting for the
>value that achieved the higher order goal at one time almost certainly
>won't do it another time. It is this latter aspect of "setting
>standards" that I don't hear reflected in your ideas about PCT
>therapy.

I have no where suggested that once someone articulates their individual standards to me that they are locked into those standards. It is the ability of people to recognize and utilize these levels to their advantage to deal with their conflicts that is important. We all change standards all the time. It is important they first recognize the existence of the standards, the part they play in how we think, their usefulness in setting and achieving references through the choices they make that is important. I don't care whether they change their standards or not. We all change standards all the time. It's being able to change within the context of avoiding or reducing conflict that is critical.

When counseling (read teaching) others, it's not what I think, it's what they think - My job is to teach them to think by helping them building confidence in their thinking ability. When they learn PCT and what goes into making up a living control system, they now have the road map. My job is to teach them how to use it. They have to learn to use it when I'm not around.

Again, Rick, I suggest you reread the role plays in Chapters 9 & 10 of Freedom From Stress. Or how I deal with the couple with whom I'm working throughout the entire book. That is how I work with others. By the way, sorry for the delay in my answer to you. Your comments have been most helpful in forcing me to rethink in this whole arena.

To Mary Powers - 90 copies of Closed Loop were printed. I made all the appropriate changes to the ever changing mailing list and I'll get an updated one to you in the mail tomorrow.

Best, Ed

Date: Sun Nov 29, 1992 5:09 pm PST
Subject: neuropsych and PCT

TO Bill or anyone,

In the past week I have been acquainting myself more thoroughly with the hierarchy as proposed by Bill, specifically attempting to figure out what species represent possession with x-number of levels, and what sorts of tasks are and are not possible given x-number of levels. Along the way I have noticed a few things which bother me and maybe you can iron out or explain. My basic observation /complaint is that we ought to be doing some more bottom-up thinking in terms of what the neurophysiological data suggest. I say this because it seems to me that we have Sensation (color) before Configuration followed by Transition (motion) when the neuropsych data suggest that color and motion are processed before "objects" are discerned (in area TE, or IT).

It seems that some levels are concerned with temporal relationships while others are concerned with spatial relationships. Some levels (Sequence) seem devoted to procedural tasks while others seem devoted to declarative tasks (Category).

OK, so we could think a little harder and come up with a different arrangement of levels, right? NO! We should take what we know about brain processing into consideration. I agree that we do not want to just go from brain data to theory--both are needed. But now it seems that the hierarchy does not jive with neuropsych data. For instance, it doesn't seem to me that processing in the parietal lobe (telling us WHERE the "object" is) is in any sort of hierarchical relationship with processing in the temporal lobe (telling us WHAT the "object" is).

How, for instance, is the Relationship level instantiated in the brain? If that is the level devoted to noting (temporal) relationships like causation and (spatial) relationships like above/below, left/right, etc, then does this level exist in multiple brain regions? Higher levels seem to reside in frontal regions.

(And one, by the way, question: if it is at the Event level that something can be perceived as a "stimulus" to an organism, does that mean that any organism on which we do research in which a "stimulus" is given is an organism with at least 5-6 levels of perceptual organization?)

Mark Olson

"It is impossible to do only one thing."

Date: Sun Nov 29, 1992 9:46 pm PST
Subject: Planning

[from Gary Cziko 921128.2330 GMT]

Bill Powers (921127.2030) in response to John Gabriel (921127.1707) said:

>The basic PCT objection to planning (as it's ordinarily done) is that
>the output of the plan is a prescription for actions that are predicted
>to have certain effects. A control system doesn't assume that a given
>action will have a particular effect. Instead, it monitors the effect,

>and is prepared at all times to alter the action if the effect deviates
>from what is desired.

And . . .

>I want to reemphasize that control systems do NOT depend on "reliable
>estimates of the effects of attempted actions on environmental
>variables." They do not attempt any estimates at all. This is one of the
>principal differences of the PCT approach from conventional ones.

I had difficulty with this PCT perspective on planning until I saw Chuck Tucker's student's videotape of a girl making brownies at the Durango meeting in 1991. Then it became very clear to me that what we normally consider planning (and a recipe is an example of a plan) is not carrying out a series of actions in order to reach a future goal but rather controlling for a series of PERCEPTIONS so that the long-term goal can be achieved. We saw the young girl repeatedly adding and dumping water to/from her measuring cup to make sure she had exactly (or very close to) one cup; mixing ingredients until she had the desired amount of blending and consistency; probing her brownies with a toothpick until it came out clean. It was very clear from Chuck's student's video that the girl was varying behavior all over the place in order to be successful in obtaining a series of perceptions that would lead to edible brownies.

Of course, just because a given plan (sequence of controlled perceptions) was successful once or twice or more in the past in obtaining some higher-level goal does not necessarily mean that it will be successful again. The yeast can be bad, the oven thermostat off, etc. But even a basic understanding of the motivation behind PCT should make it clear that repeating a series of previously successful actions will seldom if ever result in some desired consequence.

I realize that this is pretty basic stuff, but maybe it will help some people out there better understand what (I think) is behind a PCT view of planning.

--Gary

Date: Mon Nov 30, 1992 9:08 am PST
Subject: controlling perceptions? or controlling error

Bill, and anyone else,

I don't mean to make this "criticize PCT week" but I guess that's how its works out. Yesterday I wrote that we should be more concerned with neurophysiological data in developing the hierarchy. Today I want to suggest that perhaps we should change our slogan.

The following ideas are primarily those of Audra Wenzlow, a colleague of mine, so I can't take credit for them, but I do agree with them. The basic idea is this: instead of saying that we control our perceptions, we ought to be saying that we control error.

Now, lest you think that this is heretic and misses the whole point of PCT, I should mention that the concept of controlling perceptions is very well engraved into my mind. It correctly sets itself apart from traditional psychological models, and in this setting apart Audra and I have no qualms. In addition, the conceptual distinction between "controlling perceptions" and "controlling error" is not great. In controlling one, the other is necessarily controlled, assuming the control system is functional--if it isn't functional, it is the error, not the perception, which matters to the organism.

So Conceptually in terms of PCT theory, there is no problem. But there are some Important nuances that need to be addressed. First, although we can speak in terms of controlling perception (and I emphasize again that we agree with this concept), if we want to speak about "purposive-ness" then we must speak in terms of controlling error, for the organism, as an organism, has nothing INVESTED in perceptions but has a lot INVESTED in error. Perceptions are of no interest to the organism except insofar as they relate to error signals. This, I feel, is a very important point.

And this point relates to the compatibility of PCT with the rest of psychology. Isn't one of the difficulties of grasping PCT this "strange idea of controlling perceptions"--something counterintuitive since its kinda out there in the environment (almost)--controlling outputs seems so much easier to imagine since its more direct." This, I think, is what many people think--we, of course, understand "controlling perceptions." But if controlling perceptions and controlling errors are two sides of the same coin, and the latter is more easily understood (and seemingly more correct), then shouldn't we use this expression. We can say "the organism controls error and does this by controlling perceptions, not outputs." This has the added parsimonious effect of being able to say that "the organism controls error, either via a Basic Control System or via a Reorganizational Control System." This is better than saying that it controls perceptions in the former case and error in the latter.

But parsimony is not the central issue here. Neither is compatibility--we are not softening the PCT line by any means. The issue is what the organism is really doing--it is primarily controlling error, and secondarily controlling perceptions--a small distinction but an important one if we want to talk about purposiveness.

I would be curious as to why it is that we do not say it this way--is there a problem with it which we have overlooked?

Mark "It is impossible to do only one thing."

Date: Mon Nov 30, 1992 5:40 pm PST
Subject: Neurophysiology and levels; planning

[From Bill Powers (921130.0730)]

Mark Olson (921129) --

>... it seems to me that we have Sensation (color) before
>Configuration followed by Transition (motion) when the neuropsych
>data suggest that color and motion are processed before "objects"
>are discerned (in area TE, or IT).

A year or so ago we had a discussion on the net about the location of the "transition" level. There seemed to be some reasons for putting it below configurations, and some reasons for putting it above. The consensus (or whatever it was) was that both arguments seem convincing, so maybe there are two levels that do something similar. I think we agreed that there is a transition level above that of configurations specifically concerned with detecting change, but that there are also rate-sensitive processes going on below the configuration level. The lower rate-sensitive processes do not seem to give rise to a sense of motion or change perceptions of the same sort that, say, successive configurations can sustain as a sense of (stroboscopic) motion. We could certainly use some neurophysiological data on this subject.

Unfortunately, neurophysiology is too vague about most perceptions to do us much good. While the neurological observations are probably all right (identifying activity or lesions in various parts of the brain), the identification of the CORRELATES of these activities depends on informal subjective observation of the world. It isn't the "subjective" aspect I question, but the "informal" aspect.

In HPCT there are at least some principles involved in identifying correlates of brain activity. For B to be a higher level of perception than A, the existence of B must depend on the existence of A but not vice versa. When B is controlled, it must be controlled by varying A, but not vice versa. And it must be possible to control A without necessarily controlling any B, but not vice versa.

Also, to say that ANY variable is controlled, we should be able to demonstrate that human beings can maintain the variable in a constant state despite disturbances of appropriate bandwidth and amplitude. The concept of "state" includes static magnitudes, but also constant temporal patterns, relationships, categories, etc. Any really solid discussion of hierarchical relationships among the correlates of brain activity should begin by demonstrating that the proposed correlates are controllable in hierarchical fashion. That's an important kind of research that has not so far been done.

>It seems that some levels are concerned with temporal relationships
>while others are concerned with spatial relationships. Some levels
>(Sequence) seem devoted to procedural tasks while others seem
>devoted to declarative tasks (Category).

These are all categories -- yours. You're using words like "temporal" and "spatial" and "procedural" and "declarative" to name categories. So I'd classify your proposals as evidence about, or demonstrations of, category perception. Since categories are arbitrary, we can use any category terms we please for talking about any "level" (another category) we please. As long as we're just talking, anything is possible. We can say that in a first-level control system the various functions establish RELATIONSHIPS among the signals, and that when the signals change they create TRANSITIONS, and that in the overall operation of a spinal reflex there is a set of PRINCIPLES adding up to a SYSTEM CONCEPT called a control system. So we can see ALL the levels even in the first level of control systems.

What we are doing, of course, is taking different points of view within our own hierarchies and projecting them onto the system to which we're attending. It can be very hard to sort out what one is reading into some particular kind of functioning and what exists in the physical system doing the functioning. This difficulty is redoubled by the fact that we are trying to model the same kind of system we are using to do the modeling. I don't claim that I have overcome these difficulties, not by a long shot. But I think I've had more practice than most people with trying to distinguish between what is a real kind of perceptual function that ought to go into the model, and what is just one such function being applied to the activities of another.

>OK, so we could think a little harder and come up with a different
>arrangement of levels, right? NO! We should take what we know
>about brain processing into consideration.

But what we know about brain processing is based on how we organize our concepts of perceptions. All we know directly about brain processes is that different areas of the brain are active under different circumstances. The whole problem is to characterize what is different about those different circumstances. So we have to organize our

understanding of subjective perception before we can use that understanding to give meaning to activities in various parts of the brain. Without an understanding of how perceptions are related to each other, we can't understand what it means when different parts of the brain are active.

>For instance, it doesn't seem to me that processing in the parietal lobe
>(telling us WHERE the "object" is) is in any sort of hierarchical relationship
>with processing in the temporal lobe (telling us WHAT the "object" is).

Before we can say what is different about parietal and temporal lobe processing, we must be able to say what is different between "WHERE" and "WHAT." "Where" could refer to a relationship, such as "in the corner of the room." "What" could be a configuration (an elephant) or a relationship (a nail puller). The same object could be seen as a configuration or as a functional relationship. The problem is that higher-level perceptions are built out of lower-level perceptions. If the parietal lobe happened to be concerned with category perception and the temporal lobe with relationship perception (I'm not even trying to pick correct examples), then if a category is being perceived that is labelled "in", it is quite likely that a relationship is also being perceived, a particular example of the relationship we perceive directly and then categorize along with all other "similar" relationships (thereby making them "similar").

In order to sort this out, we would have to establish experimentally that parietal activity necessarily entails temporal activity, but not vice versa. If these two kinds of activity can occur independently, we would probably guess that they were of the same hierarchical level. But that doesn't mean that parietal and temporal both process categories. It could simply mean that whatever processes categories can accept information from either the parietal or the temporal lobes.

>How, for instance, is the Relationship level instantiated in the
>brain? If that is the level devoted to noting (temporal)
>relationships like causation and (spatial) relationships like
>above/below, left/right, etc, then does this level exist in
>multiple brain regions? Higher levels seem to reside in frontal regions.

One way to approach this problem would be to start with functional research about the levels, using the Test. First, do a large study that simply establishes a catalogue of perceptions that we can prove are controllable by human beings. Second, analyze the results into levels, applying the basic tests for hierarchical relationship. Third, apply disturbances guaranteed to disturb one level while leaving higher levels undisturbed, and map the corresponding brain activity.

I do have the impression that levels are distributed across different brain regions.

>(And one, by the way, question: if it is at the Event level that
>something can be perceived as a "stimulus" to an organism, does
>that mean that any organism on which we do research in which a
>"stimulus" is given is an organism with at least 5-6 levels of
>perceptual organization?)

That depends on what you're thinking of when you say the word "stimulus." If you mean something that causes perceptual signals to go through some short stereotyped space-time pattern that is perceived as a unitary thing, then by stimulus you mean what I mean by event. If, on the other hand, you're thinking of a steady pressure on the skin, or a feeling that the room is too warm, or the taste of mint, then you mean by "stimulus" what I mean by "sensation." To relate informal terms like "stimulus" to the more formal

definitions of HPCT, you must sight past the word to its meaning for you. Then you have to see how the terms of HPCT might apply to that meaning.

The only way to find out what level of perception an organism can control (regardless of what you present to the organism's senses) is to find out whether the organism can control that level of perception, given the means. You may think the organism is responding to an "event"- type stimulus, when in fact it is merely starting to control a configuration disturbance, and then relaxing when the disturbance is removed. To see if it is controlling for an event, you have to establish a task in which the organism must produce the event over and over, and change the actions by which it produces it when disturbances occur.

Gary Cziko (921128.2330 GMT) --

Thanks for the brownie example. That illustrates my point about planning perfectly, and reminds us that we CAN plan perceptions (as long as they're the kind we can reliably control).

Best to all, Bill P.

Date: Mon Nov 30, 1992 6:56 pm PST
Subject: controlling error

[From: Bruce Nevin (Mon 921130 12:12:40)]

(Mark Olson Mon, 30 Nov 1992 09:38:53) --

I like the suggestion that PCT be presented as the control of error rather than as the control of perception. It raises immediately the core question "what is error?" and spotlights the reference perception (signal). It also gets quite quickly to important questions about gain in the loop, and variations of gain that might be related to attention. More immediately, as Mark pointed out, it gets into the matter of reorganization. And it makes it more directly apparent that one must assume the point of view of the model: we are accustomed to imagining that others share our perceptions, but it seems more obvious, at least to me, that the error signal is not a public resource.

To say a control system controls (input) perception rather than behavioral output gets somewhat confusing when (one's perceptions of) one's behavioral outputs are among the perceptions that are controlled. This must be so with any conventionalized or institutionalized behavior, where one controls for conformity to norms; language is an example.

To say "behavior is the control of error" lacks the rhetorical punch of saying "behavior is the control of perception." But punchy slogans might be unwarrantedly bristly and pugnacious. This seems to me to be the sort of bridge-building that does not compromise one inch.

Unless someone sees some compromise or betrayal of principle in this that I am missing?

Bruce bn@bbn.com

Date: Mon Nov 30, 1992 6:59 pm PST
Subject: A and Beta Transfer Functions.

[Gabriel to NET 921130 12:01 CST]

I think this might be relevant so some of the issues in planning etc. Trouble is, what's new and relevant to some folks is old hat to others, and totally irrelevant to yet others. But here goes anyhow.

The overall transfer function (net gain, not OP AMP gain) is

$$G = A/(1 - A*\beta)$$

where A is in ordinary electronics, the gain of the "amplifier" and beta the "gain" of the feedback network. For large A, G is roughly 1/beta. In the first part of Bill P's book, $G = A/(1+A)$, i.e. there is an implicit $\beta=1$. It's really important to start off discussing that case because feedback theory can quickly become quite complicated.

There are some issues that don't always need to be looked at. First, in many systems, the division between A and beta subsystems is fairly arbitrary. Second, in most electronics beta is at least a constant, and its' purpose is stabilise G against variations in A. But there are lots of examples where this isn't true. They are not always called feedback systems - the Blumlein/Miller Integrator is a quite well known case.

Now, let's look at the problem of steering the EXXON VALDEZ or the Queen Mary to dock. The elements you want to control are velocity and position. You want the velocity to become zero, when the position is at the dock. The things you can actually control are engine RPM and rudder deflection, i.e. that is what the captain, through the helmsman and the engineers can control. The captain has thus to solve a differential equation for linear velocity and change of heading so as to reach zero velocity and zero acceleration, and few other rates of change at the dock.

If you put the A network between the captain's eye, and his larynx, the beta network contains the helmsman, the engineers, and a whole lot of inertia, hydrodynamics and other things.

If you are mainly concerned with the neurophysiology of reaching out for glass of water and picking it up, then the beta network in this model is simply imaging or relative positions of fingers and glass, and subsequent relative position of glass and mouth after the glass has been grasped, but before the water is drunk.

This is a very good example for neurophysiologists to look at. But if you are concerned with designing autopilots, as Martin is, then you probably want to focus on the aircraft dynamics.

If you are concerned with Government and Diplomacy, as I am at present, then you probably want to focus on the captain, the helmsman, and the engineers, and the communication between them.

Some of this can be seen in a simple model of docking the liner. To a first approximation, the Laplace Transform of beta is $1+T*p$ where T is a characteristic time needed to change anything the liner may be doing. As $T \rightarrow \infty$, $G \rightarrow 1/T*p$, our old friend the Blumlein/Miller Integrator.

And this should not surprise anybody. The whole object of the exercise is to choose any one of many possible paths to the dock, where you arrive correctly oriented, and with the last scraps of kinetic energy bled off by hydraulic friction in the last few feet.

Or, if you prefer fixed wing aircraft, a proper flare just above the runway in your landing, and enough runway left to dissipate K.E. at stall speed into the brakes.

For rotary wing aircraft, an analogous, and hairraising maneuver for those who experience it the first time, it is landing by autorotation after losing engine power.

Date: Mon Nov 30, 1992 6:59 pm PST
Subject: Re: planning

[Martin Taylor 921130 13:40] (Bill Powers 921130.0730)

>Gary Cziko (921128.2330 GMT) --
>Thanks for the brownie example. That illustrates my point about planning
>perfectly, and reminds us that we CAN plan perceptions (as long as
>they're the kind we can reliably control).

I thought Gary was introducing the brownie example to illustrate the illusory nature of planning--that the brownie cooker was not planning, but controlling perceptions on-line. I read it as Bill (I think) reads it, that the higher level (sequence or program?) for making brownies was planned and that the plan could be executed because the perceptual control at lower levels was nearly guaranteed to be possible. This high level of believability in the lower-level perceptual control gives the higher-level a predictable world in which plans are reasonable constructs.

But this makes a little dent in the idea that planning is not normal at ANY level of the hierarchy above the lowest. If the level below is actually able to control its perceptual signals, and these are fed as sensory input to the level of interest, then the level of interest is seeing a predictable world, and could execute a plan without the need for control of perception. It would KNOW what the perceptions were going to be, although not how they were going to be brought to their reference levels.

We do not think that the hierarchy works this way. Control is needed at all levels. And yet there is that element of predictability, the assurance that lower-level perceptual control is likely to work. But it cannot work perfectly all the time. If it did, the lower levels would not need to pass any information about their perceptual signals to the upper levels, because the upper levels could get the same information from the reference signals sent to the lower, information they already had. In other words, imagination would be as good as observation of the world.

What this all seems to mean is that the information the lower levels send to the higher is really information about how the actual perception deviates from the reference perception--in other words, the error. (It is actually less than that, because the higher-level probably has a pretty good idea of the dynamics of the lower one, and about what to expect about the course of error correction at the lower level). To the extent that lower-level perceptual control is good, very little information needs to be sent to higher levels. The bandwidth is low, and things normally move slowly.

I'm not suggesting that the error signal is fed upward in place of the perceptual signal; only that the information in the perceptual signal is of low bandwidth, from the viewpoint of the higher ECS, to the extent that the lower level maintains good control. (Remember that information is measured by changes in subjective probability distributions that can be seen differently from different viewpoints, and is NOT a property of a channel). The "slowing factor" can therefore be seen as a way of taking advantage of

this low information rate to reduce the channel bandwidth and thus enhance the stability of the higher-level feedback loop.

Planning, in this way of looking at the behaviour of the hierarchy, is a very effective way of enhancing the stability of the hierarchy. The problem arises with what we might call "over-planning": reduction of the information rate from lower-level perceptual signals below what is justified by the lower-level success in control. Overplanning is what causes the ill effects excoriated in criticisms of "cognitive" approaches to psychology.

Brownie points to Gary.

Martin

Date: Mon Nov 30, 1992 7:01 pm PST
Subject: Details, who needs them??

[From Ray Allis (921130.1012)
re: Other people's discussions of modelling etc. - this is from the
connectionists mailing list.]

----- Begin Included Message -----

>From ml-connectionists-request@Q.CS.CMU.EDU Sat Nov 28 05:40:04 1992
Date: Wed, 25 Nov 92 18:59:54 PST
From: Jim Bower <jbower@cns.caltech.edu>
To: connectionists@Q.CS.CMU.EDU
Subject: Details, who needs them??
Content-Length: 5091

>I grow tired of defending the validity of models to biologists
> who do not seem satisfied with any model that does not capture
>every last nuance of complexity or that does not explain every last
>experimental finding.

In response to this and several similar statements, I have to say that in now many years of building biologically motivated neural models, and attending neural modeling meetings of all sorts, I have never yet met such a neurobiologist. I have of course met many who object to the kind of "brain-hype" that originally prompted my remarks.

However, there can be no question that the issue of the level of detail necessary to account for brain function, or to do something really interesting with neural networks is a subject of active debate. With respect to neural networks, I would point out that this question has been around from the beginning of neural network research. Further, not that long ago, many believed, and argued loudly that simple networks could do everything. Several of us said that if that were true, the brain would be simple and because it is not, it is likely that artificial networks would have to get more complex to do anything real or even very interesting. As we head to the NIPS meeting, it is fairly clear that the simple neural networks have not done very well evolutionarily. Further, the derivatives are clear.

With respect to the detail necessary to understand the brain, this is also an area of active debate in the young field of computational neuroscience. However, from our own

work and that of others, maybe it is time to make the statement that it appears as though the details might matter a great deal. For example, through the interaction of realistic models and experimental work, we have recently stumbled across a regulatory mechanism in the olfactory cerebral cortex that may be involved in switching the network from a learning to a recall state. If correct, this switching mechanism serves to reduce the possibility of the corruption of new memories with old memories.

While it would be inappropriate to describe the results in this forum in detail, it turns out that the mechanism bears a resemblance to an approach used by Kohonen to avoid the same problem. Further, when the more elaborate details of the biologically derived mechanism are placed in a Kohonen associative memory, the performance of the original Kohonen net is improved. In this case, however, the connection to Kohonen's work was made only after we performed the biological experiments. This is not because we did not know Kohonen's work, but because the basic mechanism was so unbiological that it would have made little sense to specifically look for it in the network. The biological modeling now done, we can see that Kohonen's approach appears as a minimal implementation of a much more sophisticated, complicated, and apparently more effective memory regulation mechanism.

While it is not the common practice on this network or in this field to point out ones own shortcomings, it turns out that we did not know, prior to doing the realistic modeling, which biological details might matter the most. Once they were discovered, it was fairly trivial to modify an abstract model to include them. The point here is that only through paying close attention to the biological details was this mechanism discovered. From this and a few other examples in the new and growing field of computational neuroscience, it may very well be that we will actually have to pay very close attention to the structure of the nervous system if we are going to learn anything new about how machines like the brain work. I acknowledge that this may or may not be eventually relevant to neural networks and connectionism as I am yet to be convinced that these are particularly good models for whatever type of computational object the brain is. However, if there is some connection, it might be necessary to have those interested in advancing the state of artificial networks seek more information about neurobiology than they can obtain at their favorite annual neural network meeting, from a basic neurobiology textbook, or from some "overview" published by some certified leader of the field. Who knows, it might even be necessary to learn how to use an electrode.

Jim Bower

For those interested in a very general overview of the work described above, I have placed the following review article in neuroprose:

"The Modulation of Learning State in a Biological Associative Memory: An in vitro, in vivo, and in computo Study of Object Recognition in Mammalian Olfactory Cortex." James M. Bower

To retrieve from neuroprose:

```
unix> ftp cheops.cis.ohio-state.edu
Name (cheops.cis.ohio-state.edu:becker): anonymous
Password: (use your email address)
ftp> cd pub/neuroprose
ftp> get bower.ACH.asci.Z
200 PORT command successful.
150 Opening BINARY mode data connection for bower.ACH.asci.Z
```

226 Transfer complete.
bytes received in ## seconds (## Kbytes/s)
ftp> quit

----- End Included Message -----

Date: Mon Nov 30, 1992 7:07 pm PST
Subject: RE: More testing

[From Rick Marken (921130.1200)]

Tom Bourbon says:

>Who is the nut who keeps running all of those "tests?"

C'est moi!

>Hi, Rick. Does the new address work as well on receive as on send?

Yes, I am now receiving mail at the new address; but I will stay connected at the old address too for a bit, until I learn to access this LAN from home (in my efforts to revive the economy through consumerism I purchased a new computer over the Thanksgiving Holiday; my personal economy is now up s**t's creek but the California economy should be skyrocketing soon.

Hasta luego Rick

Date: Mon Nov 30, 1992 8:41 pm PST
Subject: Controlling error; Beta; planning

[From Bill Powers (921130.1300)]

Mark Olson (921130,0938) and Bruce Nevin (921130.1212) --

>... instead of saying that we control our perceptions, we ought
>to be saying that we control error.

I see a drawback in this suggestion: the only goal state for error is zero. It's important to understand that internal error is what drives behavior, but it's also important to understand that perceptions are variables, and can be controlled relative to different reference settings on different occasions. A bank teller at work finds it advisable to maintain a high reference level for honesty. But the same teller may set the reference level for honesty quite a bit lower when trying to explain to a 4-year-old what has happened to her puppy (which has just been cut in two and mangled by a freight train). Honesty is not a state of being, but a variable that can be controlled relative to high or low reference levels. In either case, the error is kept at zero. But just saying that a person maintains the honesty error at zero does not tell us what level of honesty is being maintained.

The motive of trying to make control theory less confusing to the conventionally-trained intuition is a good one. But I think that talking about controlling error will just substitute one problem for another of the same kind. People will just as happily assume that the error is out there in the public world as that the perception is. In fact you

can find lots of instances in the literature where this is exactly what is done. If you say that behavior is the control of error, you will soon find that you have to explain that the error is between a perception and a reference signal in the brain, not between, say, the objective position of a car and the objective center of its lane.

The basic problem in teaching control theory is to convince people that the critical variables involved in control are inside the brain, not outside it, and that the view we have of other people's behavior does not tell us automatically what they are doing.

John Gabriel (921130.1201) --

>The overall transfer function (net gain, not OP AMP gain) is

>

> $G = A/(1 - A*\beta)$

The transfer function that holds between the reference signal and the perceptual signal is

> $G = A*\beta/(1 - A*\beta),$

isn't it? And if you use the standard diagram I use, in which the perceptual signal is subtracted from the reference signal and all other constants are positive, the sign in the denominator has to be positive, doesn't it? Otherwise, doesn't there need to be a negative sign in the numerator for a negative feedback system?

You're correct in saying that I always use a beta of 1. The reason has to do with the conversion from physical units outside the system to NSU (nervous system units), and vice versa. This conversion is basically arbitrary. I have elected to normalize to NSU, which implies some appropriate conversion constant such that a convenient unit of measure of the input quantity (a physical quantity) corresponds to 1 NSU. For a given model, this fixes one NSU at so many impulses per second. Then we can express reference signals, error signals, and output signals in NSU, with an appropriate amplification factor and conversion factor in the output function where we translate back into physical units. This has the effect of throwing all the steady-state gain of the control system into the output function, which is heuristically very convenient. If the input is measured in feet per second, you just say that the perceptual function generates a signal of 1 NSU of perceived velocity per ft/sec. Then you scale all other signals accordingly.

An alternative is to define 1 NSU as the maximum possible perceptual signal. I did this the other day in talking about setting the reference signal to 75% of maximum, and so on. This is a useful ploy when evaluating the effects of disturbances relative to the whole control range; it helps one to see how important a particular amount of error is.

>... if you are concerned with designing autopilots, as Martin is,
>then you probably want to focus on the aircraft dynamics.

This is a somewhat confused subject, in part because there are control engineers, evidently, who think that one must compute the inverse dynamics of the load in order to provide the correct signals to the output actuators. Have you seen my Little Man model? This model provides all the correct output signals WITHOUT computing inverse dynamics. I'm beginning to wonder whether control engineers have simply forgotten about this trick of analog computing, in which the inverse-dynamic effect is created without ever explicitly calculating it.

Aircraft dynamics enter the control problem for stability reasons only. My approach, demonstrated in the Little Man program, is to start by putting the highest derivative under direct control. That system is then used by another control system controlling the next highest derivative, and so on to the state of the variable itself. I didn't invent this method: I stole it from nature. This seems to be how the "postural reflexes" work. With this method, one assures that acceleration is controlled so as to produce the demanded velocity, and velocity is controlled to produce the demanded position. In an aircraft attitude control system you need another derivative. But none of this entails the control system itself doing any calculations of aircraft dynamics. Or to put it another way, the control system IS the calculation of aircraft dynamics.

>Some of this can be seen in a simple model of docking the liner.
>To a first approximation, the Laplace Transform of beta is $1/T*p$
>where T is a characteristic time needed to change anything the liner
>may be doing. As $T \rightarrow \infty$, $G \rightarrow 1/T*p$, our old friend the
>Blumlein/Miller Integrator.

I don't remember that integrator. Could you explain it?

I would start by sensing engine thrust and putting that under control. Then I would add a second level sensing velocity, and put that under control by using the thrust control system as the output function. Finally I would sense position, and control that by using the velocity control system as the output function. If the parameters are correctly set (and the settings would not be critical), specifying a reference position will automatically move the ship to the required position, with velocity and acceleration becoming zero just as the reference position is reached. The acceleration will automatically reverse as the goal is approached. None of these effects has to be calculated in detail, coordinated, or otherwise be explicitly handled.

I guess I had better illustrate this with a model you can try. Wolfgang Zocher has worked out a rather neat program that will accept simple ASCII input specifications for any control-system setup, and run the specified model, putting whatever variables you want to monitor into an ASCII file as numbers. This program will probably be greatly elaborated, but right now it can make passing simple models back and forth very practical. Wolfgang, are you reading this? It would be a great service if you were to put your program, in C, onto Bill Silvert's file server. I will follow it up with the modified version needed for Turbo C, and a compiled version that will run on PCs. Maybe someone else could compile a version for Macs. If everyone who's interested in modeling had a runnable copy of this program, then all we would have to send back and forth are the ASCII specification files. To run a model, you just type

SIMCON SPECFILE

and it runs the setup contained in the ASCII file SPECFILE.INP and puts the output data into SPECFILE.DAT, another ASCII file. Later we'll work out reading programs that can convert the ASCII files to graphical plots, but even as the program stands it is highly usable.

Let's have some comments from everyone out there -- are you interested in getting a copy of the program? Can you compile it from C source code (in which case it will run on just about anything)? Talking about models without running them is like explaining an aircraft maneuver with your hands in your pockets.

Martin Taylor (921130.1340) --\

Some very well-made points about planning. If you can plan, you don't need to control. I would add that at any level, there are uncontrolled perceptions that can introduce disturbances, so that even if all controlled perceptions at a given level are perfectly controlled, the higher levels still work with partially-uncontrolled inputs.

>I'm not suggesting that the error signal is fed upward in place of
>the perceptual signal; ...

That, however, is one viable form of the model, probably applicable mainly to higher levels. Maybe if Wolfgang gets his program into wide use here, we can pass some models around and explore this alternativemodel-based form of the control model. It has some interesting possibilities, including the ability to continue controlling with intermittent input data. Such a system is in the imagination mode ALL THE TIME, but there is provision for continually comparing the model's behavior with real perceptions and using the difference to modify the model. This requires error signals from lower level systems to be passed upward, not perceptual signals.

Best to all, Bill P.

Date: Mon Nov 30, 1992 8:49 pm PST
Subject: controlling error

Let me just say one thing, concerning Bruce's post: one might want to say that "behavior controls perceptions" yet "the organism controls error." Possibly a meaningless semantic quibble--possibly not. This goes along with what Bruce was saying concerning one's perspective. Saying that behavior controls perceptions is a valid observer approach, while saying that the organism controls error takes an individual system's approach--Bruce said it better. Possibly, each is equally good within its own approach--if this is true then I would contend that the individual system's approach is better.

But perhaps I shouldn't have posted this for it may swing the discussion into semantic quibbles which is not what my initial post was about.

Mark "It is impossible to do only one thing."

Date: Mon Nov 30, 1992 9:10 pm PST
Subject: A little bit of background

From Bill Cunningham (911129.1600)

It's taken most of today to relieve the constipated queue on the Ft Monroe mailer, noting that short term results aren't always desirable. The ongoing Powers/Gabriel/Taylor discussion is an outgrowth of a project (great quest) that John, Martin, Tom Baines, a non-netter, and I have undertaken. We are preparing a 4-hour tutorial (forceably reduced from 7 hours) entitled "Optimal Decision Support: A Perception Based Paradigm for Decision Centered Information Fusion." Although a mouthful, the title is intended to focus attention on what's wrong with the current theology for military command & control (C2)--and offer some wisdom we deem essential.

Our topic really has no unique military association, but that serves as a handy forum. Our audience will be a closed symposium sponsored by the Joint Directors of (US DOD) Laboratories, called the Data Fusion Symposium (DFS). The problem we address is how an organization collects information from a myriad of disparate sources, processes

(combines, folds, spindles, mutilates) the information into the basis for effective decisions by senior executives in the organization. Bluntly put, we know how to generate, store and transport more information than we know what to do with. The executive is buried under a glut of useless information and may even seek more in an attempt to get the "right" information.

The current process is driven by the purveyors of fine information, rather than by the consumers. The term "data fusion" refers to various aspects of correlating and combining input data. The term "information fusion" was coined to make the point that we propose something much more than a better way to do the same old thing.

The makers of fine decisions see that as their primary function. They are beginning to articulate their frustration with the current process. That's new. They used to just complain. We find we are right on track with the customer.

We would not be having this discussion if it were not for Bill Powers and PCT. The fact that "data fusion" can never really address the problem isn't apparant until you realize that the executive can only act on his/her perceptions, and that the so-called decisions are really behavior aimed at producing the desired percept.

In response to Bill P's comment about designing an organization based on PCT, I would respond that we (at least me) are merely observing what really goes on. This whole business started two startling realizations: (1) the PCT hierarchy maps directly into a hierarchy of aggregation/abstraction within the organization; (2) when the organization is through processing the information, it must be transferred into a real live individual control system. This led to the observation that humans seem to have created organizational functions that immitate what they do internally. That may not be universally true, but it seems to fit here.

We have been focusing on the organizational sensory chain, not designing an organization. The comment that an organizational hierarchy should specify the desired percept to the subordinate echelons is spot on. We just haven't gotten there yet.

As much as it may surprise some readers, I feel the Army is light years ahead of the other services with respect to PCT (at least in the US). They don't call it that, of course. But the guiding philosophers have seen the need for distributed control--since the force is spread all over creation and each local commander must respond to his/her local world as it is perceived. So I hear very senior officers saying things like "We can't all have the same view of the battlefield because we each see from a different perspective. But I want us to hold in common those things that matter. I want my subordinates to understand my basic intent and guidance and take it from there." "I want to see _____ result. Make it happen." "I want to control the filtering of information coming to me. I don't want to be the slave of what somebody else thinks I need to hear." "What's signal for me may be noise for you and vice versa. I want to control my own signal-to-noise." The latter is a direct quote from our current 4-star commander who has thought the problem through and has borrowed terms he's heard to describe what he wants. The quote is rather funny because it came out of a spirited discussion with an info purveyor who kept telling him he couldn't control the flow without being inundated by the volume. The discussion ended when one of the two (your guess) said "I obviously haven't raised my signal above your noise level."

Bottom line is that we have a pretty solid outline of what needs doing and PCT is at the center. But there are other aspects of the problem I'll write about later. Interestingly enough, our audience is neutral on PCT and has no axe to grind there. But our message

that the current approach to data fusion is bankrupt sure will upset some folks. We live in interesting times.

Best to all,

Bill C.

Date: Mon Nov 30, 1992 11:09 pm PST
Subject: Misc remarks; A and Beta Transfer Functions.

[From Rick Marken (921130.2000)]

Ed Ford --

Your last post on standards cleared up a lot. I know that what you do is teach people to control their own lives more effectively; sometimes, I take issue with the way you describe some of your therapeutic goals. But your last description was excellent and quite compatible with my own sentiments about therapy (and they are just sentiments, since I would never be able to actually DO therapy as skillfully as you do it.)

Mark Olson --

I'm not sure I like the idea that it is the error signal that is controlled. To me, control means producing a result that matches a specified standard. There is really no specified standard for error -- the control process (when it works) does tend to move error toward zero, so there is an implicit standard (zero) for ALL error signals. But this, I think, misses one of the most important aspects of the control exhibited by living systems; controlled results are made to match standards that often vary over time. So it is these results (perceptual variables) that are controlled, relative to fixed or varying specifications for standard. If we watched the variability of an error signal over time, it's behavior would be decidedly uninteresting; it would just wobble around a bit. If, however, we watched a perceptual signal and its reference specification, the behavior would be far more interesting -- the perception would track rather precisely the any variations in the value of the reference; and this would happen despite visible disturbances that should keep the perceptual signal from tracking the reference in this way; much more interesting behavior; much more powerful observation.

So I vote for sticking with the idea that it is perceptual variables that are controlled -- for the sake of both rigor and aesthetic satisfaction (once you develop the imagery to support the mathematics of the theory of control, it really is quite beautiful).

Gabriel to NET 921130 12:01 CST--

>Now, let's look at the problem of steering the EXXON VALDEZ or the
>Queen Mary to dock. The elements you want to control are velocity
>and position. You want the velocity to become zero, when the position
>is at the dock. The things you can actually control are engine RPM
>and rudder deflection, i.e. that is what the captain, through the
>helmsman and the engineers can control.

Bill. I think you got some 'splainin' to do.

I would say that what the captain can "actually" control is only his/her perception of the velocity and position of the boat; engine RPM and rudder deflection are among the

variables that influence these perceptions and allow control. The rudder is not actually controlled (since it is not sensed) but some sensed indication of rudder position may be controlled in order to control position and velocity of the boat.

>The captain has thus to solve
>a differential equation for linear velocity and change of heading
>so as to reach zero velocity and zero acceleration, and few other
>rates of change at the dock.

This is precisely unnecessary (fortunately so, since it is also impossible to find a solution that works in the continuously changing environment of the vessel) because the captain is controlling only the perceptual results of his/her actions.

Behavior (fortunately for those of us who have a hell of a time solving differential equations) is the control of perceptual input, not behavioral output.

Best Rick M.