9202 CSGnet

Date: Sat Feb 01, 1992 6:08 am PST
Subject: CLOSED LOOP

From Greg Williams (920201)

>[Gary Cziko 920131.2025]

Thanks for your kind comments about the latest issue of CLOSED LOOP.

>Have you guys thought of single-issue distribution [of CLOSED LOOP], or would
>this be too much of a hassle?

Ed Ford is handling ALL aspects of distribution (thank goodness!), so he's the one to decide. Just as long as it makes lots of profits for CSG, I'm up for additional distribution. Even if it doesn't!

>P.S. The only quibble I have with _Closed Loop_ is the spelling of CSGnet.
> Greg writes it "CSGNet." I like the small "n" so that the two meaningful
>units of CSGnet are nicely distinguished. Since I started CSGnet, I want
>it spelled MY way. What do you think of that?

Gary, you should have spoken up sooner; I thought I was making an editorial improvement! In future issues, not only can I use "CSGnet," I can even put in a copyright notice if you want: "CSGnet name Copyright 1990 by Gary Cziko." You'll have to see about getting a trademark on it yourself, though. And maybe on last year's issues and Volume 2, Number 2, you can take little pieces of paper with "CSGnet" on them and paste them over each occurrence of "CSGNet" -- may subscribers send back issues to you for retrofitting? (Just kidding.)

"CSGnet" it is!!!

Greg (co-publisher of HortIdeas newletter, PictureThis and AniThis software, which maybe should have been, respectively, Hortideas, Picturethis, and Anithis?)

P.S. Our reply about BEER'S BUG is coming along -- we were busy yesterday making a new composting privy. The old one lasted about 14 years.

Date: Sat Feb 01, 1992 8:47 am PST
Subject: Re: CLOSED LOOP

[from Gary Cziko 920201.1000]

Greg 920201

I said:
>>Since I started CSGnet, I want
>>it spelled MY way. What do you think of that?

You replied.
>"CSGnet" it is!!!

Sociologists take note. Another demonstration of the power of social influence!

>Greg (co-publisher of HortIdeas newletter, PictureThis and AniThis software,
>which maybe should have been, respectively, Hortideas, Picturethis, and
>Anithis?)

Not at all. Since you are separating two words, it makes sense to start the second one with a capital letter. But since CSGnet is made up of an acronym and a word (really an abbreviation), we capitalize each letter of the acronym (like USA and USSR) and then switch to lowercase to show the start of the new word. I'm sure you'll eventually admit that it looks much better this way!

>P.S. Our reply about BEER'S BUG is coming along -- we were busy yesterday
>making a new composting privy. The old one lasted about 14 years.

First things first, I suppose (I realize that you can only hold some things for so long). I'm looking forward to your reply.--Gary

Date: Sat Feb 01, 1992 9:07 am PST
Subject: Closed Loop

from Ed Ford (920201.10:00)

Concerning Closed Loop: Greg Williams deserves a great thank you from the membership for the job of taking three months of CSGnet and reducing it to 72 pages. His task is an overwhelming job with little financial compensation. As far as I'm concerned, producing various issues according to subject matter is way beyond our resources, both human and financial.

To reduce our costs, we recently reduced our mailing list which now just includes 42 paid members of the CSG plus 35 people who paid their dues last year but not this year. The cost of producing and mailing Closed Loop has to be kept within our present income and we are doing just that, but with no margin to spare. We do print a few extra copies of Closed Loop for those who join the CSG during the year and for special requests, but not as many as we did with prior issues.

Ed Ford ATEDF@ASUVM.INRE.ASU.EDU
10209 N. 56th St., Scottsdale, Arizona 85253 Ph.602 991-4860

Date: Sat Feb 01, 1992 11:56 am PST
Subject: Language & perception; BEER bug

[From Bill Powers (920201.0900)]

Bruce Nevin (920130.1633) --

In your illustrations of the method of Harris, certain phrases come up that are considered by many people, perhaps even every person who speaks English at all well, as unacceptable. Can you think of any *a priori* reason why people would consider the following phrases unacceptable, even if they weren't being compared with other phrases?

The slipper chewed the puppy
The ball ate the ice cream
The ice cream ate the ball

I believe that Harris would say these phrases are unacceptable because in each case the subject of the phrase is an operator that does not take an argument, and so can't be connected by a transitive verb to an object. That isn't an a priori reason, but a generalization derived from considering many comparisons of phrases by individuals. The explanation should be stated more empirically: it has been found that ...

The *a priori* reason I'm thinking of is simply that the implied events, which we can imagine, never actually happen. We have never seen anything called a slipper, a ball, or ice cream doing the sort of thing that the verbs imply. It's hard to imagine how they could. It isn't language, but the experienced world that tells us these sentences are unacceptable. Harris' method simply asks people, "In your opinion, do the perceptions that are the meanings of these sentences actually happen in the real world?" (Or perhaps "In a world you can imagine?") By this method, Harris is not really exploring the world of language. He is exploring the world of perception.

If you were to make up another set of sentences relating to aspects of the world of which a person has had no experience, judgements of acceptability would mean nothing:

The current heated the resistor	The resistor heated the current
The voltage heated the resistor	The resistor was heated by the voltage
The current dissipated energy	The current was dissipated by the energy
The field deflected the charge	The charge deflected the field

If you don't know what the nouns MEAN -- that is, if you have never experienced anything nonverbal to go with them -- and if you have no idea whether one thing designated by a noun can do the action described by the verb to the other thing designated by the other noun, you won't be able to judge the acceptability of these sentences one way or the other. There is nothing about hooking the words themselves together in these ways that is acceptable or unacceptable. All that is accepted or rejected is what the words mean.

So Harris is not really asking his subjects how words work in these sentences. He's asking them how the world described by the sentences works. He could show animations illustrating the meanings of the sentences, and ask "Is this something you would expect to see in real life?" and you would get the same rankings for these sentences. Asking whether an operator takes an argument is asking if the thing designated by the operator is related to the thing designated by the argument in a certain way. I repeat: this is just asking how the world of perception works, with the words being used merely to indicate the perceptions.

Ambiguities in the conclusions can come from not thinking of the right perceptual context, as in the second set:

>(2)   John hit the ball   The ball hit John

These are both perfectly acceptable. John is retired, because a player hit by a batted ball is OUT. They are not acceptable if the image you get from "the ball hit John" is one of a ball that of its own volition jumps up and strikes John -- that doesn't happen, so the sentence is judged unacceptable.

I'm making certain assumptions about Harris' method here. I'm assuming that these investigations end up with lists of words or phrases that have been classified as operators and arguments according to the combinations that actually appear in language usage. "Slipper" is a word that does not take any arguments, although it can be an argument. An word like "hit" does take an argument but unless used as a noun can't be an argument. I'm just guessing.

At any rate, I assume that the outcome is an encyclopedia of words classified according to their partial orderings in sentences as they actually occur. These partial orderings Harris takes as properties of language itself. If the same partial orderings are found in another language, Harris assumes that they indicate invariants of language that are characteristic of human Language. This being an empirical study, the question as to WHY these orderings appear is not asked. That is, Harris does not look to any other level of explanation of the orderings as a way of predicting whether one ordering will be observed while another will not.

What I'm trying to get at here is that meanings explain the orderings. If you just assume that subjects in the experiments were translating from the words into imagined experiences, and then judging whether the imagined experiences can actually occur, you should get the same partial orderings that are observed, given the context of experiences that the subject has in mind (this can be determined by asking). To the extent that nonverbal perceptions are the same in all human beings regardless of the language they speak, the partial orderings found by Harris should be the same in all languages. In any culture in which there are puppies and slippers, slippers are not capable of chewing puppies.

If there are differences, they may be due to culture -- perhaps in one culture "the slave sold the master" would be unthinkable, and never happens, and would be judged an unacceptable form, even though the person can reluctantly imagine it happening and can easily imagine the counterpart of "The master sold the slave" and "The slave painted the steps." These judgments have nothing to do with language.

You can say "The puppy chewed the slipper" and "The slipper warmed the foot" but not "the foot warmed the puppy" or "the slipper chewed the foot." All the sentences are perfectly good sentences. It's just that some of them refer to perceptions that, apparently, seldom or never occur -- they don't belong in the world as we experience it. Sentences can suggest assembling imagined perceptions in certain ways. But the perceptions are boss when you ask "Is this likely?".

Avery Andrews (920130.1830) --

> ... even though we seek uniform underlying principles, there is going
>to be extensive variation in what is actually observable.

Behaviorists like to say the same thing: it isn't that our model is wrong, it's that behavior is inherently variable. I think that when you do finally arrive at fundamental principles of linguistics, the behavior they describe won't be signficiantly variable. The principles just have to be about the right things -- the things that are actually uniform across people. We can deduce that disturbances of controlled variables will result in equal and opposite efforts by the controlling system, and this is an extremely uniform finding. But if we tried to say what disturbances will occur, and exactly what behaviors will be used as a means of opposing their effects, we would find extreme variability. We can say that all people perceive a type of variable called a configuration, and that would be true of essentially all people and (I guess) all cultures. But if we tried to say that everyone perceives apples and squares, the variability would show up again because people are not uniform at that level of detail.

I think that when you view language as part of a process of controlling perceptions, regularities will show up that are not apparent in an analysis of usage that deals with language just as manipulations of words. To me, variability in data is a sign that we're looking at the wrong data, or looking at it from the wrong point of view.

Martin Taylor (920130.1905) --

I think you've clarified matters concerning repetition. There's a continuum, now that you mention it, extending from completely irregular spacing to completely regular spacing (or timing). A "repetitiveness" detector might be a Fourier transform yielding a single spike of maximum amplitude for the regular spacing (simpler mechanisms exist, however). As randomness appears in the spacing, the amplitude of the response becomes less and the tuning becomes broader. When there is just a random assortment in the field, the amplitude becomes zero and the broadness of the response maximum.

Repetitiveness is independent of the sense of what is repeating:

```
    T  T  T  T  T  T  T  T  T  T  T  T  T  T  T  T  T  T  T  T  T

   TT     T T   T      T  T TT   T      T      TTT     T  TT   T  T
```

Both of these strings are groups of Ts, so there is a sense of many Ts. But only one gives a sense of repetitiveness that's the same over the group. I would conclude that repetitiveness and manyness are independent dimensions of perception: each can vary without affecting the other (much).

> ... introspection is a suspect way of studying what goes on in one's
>head, ...

I don't know how else you could find out what's going on in your head by way of perceptions.

> ... but I THINK that when I do symbol manipulation exercises such as
>mathematical, I use the imagic mode to do it.

Well, do you or don't you? Nobody else is in a position to know. I accept your report as a report what what you experience.

>I cannot imagine how one simply does symbol manipulation as such,
>without these images.

Sure you can -- try another context. How about long division or calculating square roots or converting a logical expression to minterm form? When I set up the equations for a control system in algebraic form, and solve them simultaneously, I don't bother to visualize the significance of every transformation -- I just do the manipulations according to the rules and concentrate on avoiding mistakes. There's a certain amount of imagery involved, as in "transposing" and "cancelling" and "reducing," but I don't try to get creative about it. There's some imagery in the background as I look at each partial result and muse about relationships among the underlying variables that show up. But basically, if I want to get the right answer, I just turn the crank.

I think differences among people in imaging are mainly a matter of perceptual level. I don't imagine at the sensation level except when dreaming (although I did when very young), so my mental images aren't "vivid" or "solid." I can imagine THAT I'm touching something, but I don't imagine the touch. It's the relationship signal that I experience, I suppose -- certainly not the sensation signal.

RE: money. I've wanted to get an economics discussion going. My father's done some very interesting stuff in macroeconomics (not control theory -- the old curmudgeon). I think control theory puts a missing element into macroeconomic concepts, the battery that keeps the circular flow going. If you're game, why don't we each post a short essay and see if something develops?

Gary Cziko (920131.0931) --

RE: Beer's cockroach model

Pretty sharp, Gary:

>Why would "those little feet fly" if leg movement was feedforward?

If the legs move exactly the same no matter what the footing, then the movements are strictly feedforward. But if the scrabbling speeds up as the feet slip, you'd conclude at least that forward velocity is the controlled variable.

Feedforward is really a terrible term: all it means is an input-output process. In spinal reflexes, muscle contraction is "feedforward" because the contractile fibers shorten in accord with the neural signals entering the muscle, without direct feedback to alter those signals. If you consider a larger loop, this "feedforward" function is seen as the output part of a control system, a spinal reflex involving feedback from the consequences of muscle contraction.

Leg POSITION in the cockroach could possibly be passively dependent on outputs of higher systems. With a signal the leg snaps forward; without it it snaps backward. If that's the case, we just say this whole action is part of the output function. Control theory doesn't contain any premises about what variables are controlled. You have to find out what the controlled variables are by observing the real system. In fact, you don't have to assume you're looking at a control system: just analyze the relationships you can see, and find out what kind of system it is. If it turns out to involve negative feedback and reasonable loop gain, well, that's what we call a control system.

I think Greg is right about obstacles constituting disturbances. The model makes the cockroach turn when its antenna touches the obstacle; the turning affects the receptor signal. The reference level is evidently zero antenna signal (a zero reference signal doesn't require an explicit reference signal to exist). If you move the obstacle toward the cockroach and touch the antenna, the cockroach will veer away from it. Perfectly good control system, if a little crude.

Date: Sun Feb 02, 1992 9:42 am PST
Subject: BEER'S BUG

FROM PAT & GREG WILLIAMS (920202)

>Gary Cziko 920131.1000

>Why would "those little feet fly" if leg movement was feedforward?

On page 74 of Beer's book is the cockroach locomotion circuit found by neurophysiologists. The gait rhythm is controlled by a central oscillator which is modulated by sensory inputs which occur if a leg is far enough forward or far enough backward. What is put out to the muscles is a force signal. If a leg is in the middle of its possible positions (neither far forward or backward) and the foot is down, then the speed of movement of the leg (backward) relative to the body (which is actually moving forward relative to the ground) is determined by the "feedforward" force output. If the feet are not slipping on the ground, the legs will tend to move slower than if the feet are slipping, because the physical forces opposing the forward movement of the bug are greater in the former case (the bug will move faster relative to the ground). With feet slipping and therefore faster leg movement relative to the body, the backward angle sensors will be activated with a smaller cycle time -- and those little feet will fly. In Beer's model, he didn't really put in all of the F=ma physics.

>If we could show that a cockroach maintains its walking speed
>over surfaces of varying slipperyness by varying the temp of its gait then
>we have the classic demonstration of a controlled variable. It seems to me
>that the little feet should NOT fly if it is feedforward.

The gait speed goes up with slipperiness not because of control circuits inside the bug, but because, as we said above, the force to the muscles remains the same (at a given point in the cycle) but the forces opposing bug forward motion are less with slipperiness (because the feet slip, and the bug doesn't move forward as fast).

>Also, doesn't flypaper work because the little critters are just not strong
>enough to pull free?

We suppose so.

>Haven't I seen flies and other bugs leave a limb or two behind when only a
>limb or two makes contact with the sticky stuff? If so, wouldn't this
>indicate that they are pulling a lot harder than normal and also evidence for
>negative feedback?

If one foot were stuck and the other five weren't, the unstuck legs probably could generate the force required to tear off the stuck leg. No "extra" pulling required (no control required).

>I had thought about how weighing so little makes one less prone to gravity
>and inertial disturbances. I suppose it doesn't make much difference if
>you're walking uphill or down if you weigh next to nothing to begin with.
>But aren't there other disturbances that don't depend on weight (we've
>already mentioned stickiness and slipperyness)?

Yes, but what if they just don't matter very much. Several invertebrates have "escape" (demonstrably feedforward) circuits which, in "reaction" to ominous sensory input, get the animals the hell away. ANY WHICH WAY away! The neural signals sent to the muscles are NOT fine-tuned by feedback control for trajectory -- they just say CONTRACT FAST. And the would-be prey is suddenly over THERE, or THERE, or THERE -- it doesn't matter which.

>So maybe you're saying that it is negative feedback at upper levels but
>feedforward at lower ones.?

The high-level negative feedback loops extend all the way through the environment VIA low-level feedback AND feedforward loops. The feedforward loops are compensated for eventually (and, basically, periodically), but not necessarily continuously. Think of the "bang-bang" (on-off, discrete) control of a typical thermostat.

>If the neurosciences could easily find control loops wouldn't all the
>neuroscientists interestd in mammals be on CSGnet by now? It is really so
>easy to find this evidence looking at the neuroanatomy? I'd feel better
>basing my model on behavioral data.

Neuroethology COMBINES physiological and behavioral data. And, yep, there is good behavioral data for feedforward circuits in certain invertebrates, including the cockroach.

>Rick Marken (920131)

>The same is true of Beer's bugs. The disgrams he shows make them look like
>S-R systems but they are really feedback control systems because their
>outputs influence their inputs; his differential equations for the neural
>dynamics take care of the dynamic constrainst that are needed to keep the
>control system stable.

We agree, except that we don't think his diagrams make the bug look like an S-R organism.

>So Beer's Bugs already are control systems -- controlling input. And to the
>extent that those inputs are disturbed the disturbances will be resisted.

Yes.

>Everything would be clearer if Beer designed his bug from the bug's per-
>spective. Then he would be a real PCTer. As it sits, his stuff "works"
>inasmuch as it entertains observers but it doesn't provide any fundemental
>understanding of how behavior MUST be organized (around the control of
>perception) if it is to be ADAPTIVE (disturbance resistant).
>So my main complaint is that the Beer (and Brooks) approach obscures what
>is important. It would be very hard to build a fairly complex control
>system using their architecture (one that controlled many, complex input
>variables) though I agree that it could probably be done. But I think it
>would be like trying to do arithmetic with Roman Numerals.

Nobody said it would be easy (or "cognitively penetrable"!), Rick. PCT's "functional-level" models appear easier to grasp, but all of the nitty-gritty details will have to be supplied to make them work in the real world. Just ask Bill Powers about his having to POSTULATE the recognition mechanism for visual objects of interest to his Little Man. If the Little Man were built as a robot, that mechanism would have to be built, not just postulated (and handled by computer software constructs). The problem with high-level models (as traditional AIers have found out) is that when your organism and environment are entirely conceptual (software), handling the low-level stuff is TOO EASY because it is solely representational (in the CP sense). But what can one build out of fairly realistically modelled neurons only? Beer's bug, for a start. (We cheat a bit here on this claim, since Beer also modeled the sensory inputs representationally -- but he DIDN'T (as Chris Malcolm will attest) model the bug's nervous system representationally). This is the BIG DEAL of Beer's bug.

We're working on an expanded Nervous System Construction Kit which will allow PCTers to apply their ideas to building nervous system models, and we hope they will do so.

>Bill Powers (920201.0900)

>If the legs move exactly the same no matter what the footing, then the
>movements are strictly feedforward. But if the scrabbling speeds up as
>the feet slip, you'd conclude at least that forward velocity is the
>controlled variable.

See above.

>Feedforward is really a terrible term: all it means is an input-output
>process. In spinal reflexes, muscle contraction is "feedforward" because
>the contractile fibers shorten in accord with the neural signals entering
>the muscle, without direct feedback to alter those signals. If you
>consider a larger loop, this "feedforward" function is seen as the output
>part of a control system, a spinal reflex involving feedback from the
>consequences of muscle contraction.

Good point. Local "feedforward" can be embedded in global feedback loops. But that doesn't make it local feedback!

Best

Pat & Greg

Date: Sun Feb 02, 1992 1:05 pm PST
Subject: Emotions

[From Kent McClelland 920202]

Bill Powers,

I see from the latest copy of the newsletter that you're putting out a new collection of essays, LIVING CONTROL SYSTEMS, V. 2, due sometime this spring. The announcemnt notes that this volume may include, among other things, "a 'missing' chapter from BEHAVIOR: THE CONTROL OF PERCEPTIONS on emotions. . ." This reminded me that I had promised my seminar class that I would ask you on the net to comment on how emotions are generated in the control theory model and what part, if any, they play in maintaining control of perceptions or interfering with it. The class have been briefly introduced to control theory and are still, of course, trying to get some notion of what it's all about, but a couple of the bright seniors in the class offered the criticism that PCT, in common with several of the other theories they have looked at, seems to have relatively little to say about emotions. They're of the opinion that understanding emotions is a pretty important part of understanding human behavior.

My request is, then, that, while we're waiting for the book to come out, you give us a capsule summary of some of the things the "missing" chapter has to say about emotions. Are emotions simply epiphenomena of control or lack of it, or do you see them as playing a more significant role in the hierarchical model?

Dag Forssell

You asked sometime in the last couple of weeks for the names of people conducting courses that deal with control theory. As you can see, my name should be added to the list. I'm teaching a seminar this spring that takes control theory as a starting point for discussion of theories and viewpoints on social control.

Best to all,

Kent

Kent McClelland   Office: 515-269-3134
Assoc. Prof. of Sociology   Home: 515-236-7002
Grinnell College   Bitnet: mcclel@grin1
Grinnell, IA 50112-0810

Date: Mon Feb 03, 1992 4:33 am PST
Subject: from neural-net discussion

[From: Bruce Nevin (Mon 92013 07:22:18)]

Kanecki's description of the SA article caught my eye. This and other aspects may be of interest, discard if not.

Bruce
bn@bbn.com

-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-

Neuron Digest   Wednesday, 29 Jan 1992
                Volume 9 : Issue 3

Subject: Having a machine perform a simple task: The reason why neural networks and artificial intelligence need each other
From: David Kanecki <kanecki@cs.uwp.edu>
Date: Sun, 12 Jan 92 17:11:30 -0600

In the december issue of Scientific American an article described an attempt to build a robot to butter a piece of toast. In the article this task was not as easy as it seems. For example, if the toast was moved to a different spot a new set of rules and procedures were needed. Also, if the toast was oddly shaped a new set of rules and procedures were needed. The solution to the problem was to build a mini netowrk to adapt to changes.

Rules as used in artificial intelligence work in task that are fixed and non changing. And, neural networks are usefull in changing environments. In the natural world, each task may follow a set of rules. Also, the rules may need to altered slightly due to the environment. Thus, for a machine to function in the natural world it needs to have an artificial intelligence and neural network basis. One movie that showed this connection was "Star Wars". In "Star Wars" there were two robots. One robot was named r2d2 and the other was c3p0. The r2d2 robot was the artificial intelligence basis, while the c3p0 was the neural network basis. Lastly, in complex task both had to work together.

The development of self learning games is one way to develop a program that utilizes an artiificial intelligence and neural network basis. In my experience, learning and task based program still cannot breakaway from the Lady Lovelace argument. -- "A machine can only do what it is programmed to do" or the neural network corrollary - "A machine can only learn what it is taught by a teacher".

Thus, the goal of my study is to try to develop a program that can teach and solve problems itself by using techniques of ai, neural networks, and biology.

David H. Kanecki

kanecki@cs.uwp.edu

===

Subject: DuPont Neural Computation Program - Job Opening
From: elaine@central.cis.upenn.edu (Elaine Benedetto)
Organization: University of Pennsylvania
Date: 15 Jan 92 16:05:27 +0000

Job Opening in
Modelling/Simulation of Neural Systems
****

The DuPont Neural Computation Program invites applications for a computational modelling and simulation position. Applicants shuold have experience in model- ling techniques and be interested in general problems of sensory/motor integration and/or single neuron computation. The program constructs biophysi- cal model neurons reflecting experimentally recorded neuronal dynamics, and then assembles these neurons into networks reflecting the organization of a biological cardio-respiratory control system. Familiarity with the UNIX/C/C++/ X-windows computing environment is desirable. Interaction with ongoing experi- mental work will be required. We are open to considering various levels of academic training and work achievement, but imagine that the ideal candidate might hold a B.S. in a technical subject while having exposure or interest in biological subjects such as biophysics and neurobiology. The opening is immediate, but a later start date will be considered. Competitive salary commensurate with experience.

Contact: Dr. James Schwaber or Prof. Lyle Ungar
Email: schwaber@eplrx7.es.duPont.com ungar@cis.upenn.edu
Telephone: (302) 695-7136 (215) 898-7449

DuPont Neural Computation Program
DuPont Experimental Station E352 Room 253
Wilmington, DE 19880-0352
==

Date: Mon Feb 03, 1992 2:19 pm PST
Subject: language learning

[from Joel Judd]

(I've had no access to CSG since Thursday, so I apologize if any of this has been discussed since then)

Rick & Martin (920130),

I appreciate the re-summary of learning. It was helpful to be reminded of the intricacies of hierarchical change--I had been lumping all change together into sort of a unitary event. I'd like to repeat back the three types of change with examples to check my understanding. I'll use the old standby, phonology.

------- I--------

>1) a change in the characteristics of the sensory function that maps inputs into perceptual signal; >such reorganization changes the "meaning" of the perceptual signal.

I assume this means that given the same environment, the only thing changing is how I percieve it. Is this as simple as old psych visuals like a Necker cube or the old lady/young woman?

>2) a change in the output function which maps error signal into lower order outputs...

Often, speakers from languages with a "flap r" (eg. Spanish) don't use it when saying words like 'butter,' where the double 'tt' is essentially identical to the single 'r' in words like 'pero.' But as soon as it's pointed out to them, there's no problem.

>3) a change in the connection of the system to other systems; this would include or remove the >system from other control organizations.

In thinking about the phonology example, I have a difficult time seeing it being "learned" without all three types of change occuring. For years speech therapists have recognized that "perception precedes production"; the recent mention of the Silent Way of Brown seemed to find most people in agreement. Then the person must be able to produce an articulation matching the percetpual reference. Then (speaking of an L2) it must be included in a different configuration of ECSs. I like the above example also because it explains how an existing and functioning ECS can be employed in an L2 (the old problem of language transfer).

Is the "flap r" learned? Well, the Spanish speaker already knows and uses it (in Spanish), but wasn't perceiving it in either hearing or using words like "batter." So is it a case of simply needing to PERCEIVE it in the neural context utilized by the System Concept "English" (change #3)? Are other changes (1&2) reserved then to explain novel perceptions (eg. a Spanish speaker learning a "retroflex r")?

------------ II ------------

I have usually seen "Intrinsic Error" used in relation to physiological needs of the system. Yet learning is often discussed regarding activities that don't seem to be "life or death." Does (most) learning become less "critical" the older we get (especially after the first 3-4 years), and is that a reason why there is so much variation (really now, per the example above, who gives a @&%# whether or not I learn the "flap r" in English)? What does it really mean to say that all reorganization results from Intrinsic Error?

--------------- III --------------

In connection with Rick's comments on randomness, it has also been suggested and promoted by Gary that Blind Variation and Selective Retention (BVSR) process outlined by Campbell seems to be a promising description for the reorganization process

Bruce and Bill (920130)

>To the extent that language is conventional, its structure is optional. Only those aspects of it that >are exactly the same for every single human being can reflect the true properties of and basic >functions of the brain.

So...in a very real (or very trivial?) sense, NO PART OF LANGUAGE is the same for any two people. What's left across people? Whatever aspects of perceptual control are MANIFESTED IN language. Once again, if we focus just on language, we're left with the problem of making sense of outputs, albeit very sophisticated outputs. One must posit the internal processes responsible for linguistic behavior. If we control linguistic Relationships, we control Relationships. If we control linguistic Categories, we control Categories. You know, this isn't gonna sit well with some language researchers. And I love it.

Command: prinht t 1-3

Date: Mon Feb 03, 1992 7:09 pm PST
Subject: Re: learning

[From Rick Marken (920203b)]

Martin Taylor (920131 19:30) says:

> What, for example, is "intrinsic error"?

Good question. I think I should start off by saying that my earlier posting on reorganization was a description of how it works IN THE MODEL. So the term "intrinsic error" refers to something in the model, viz, difference between reference and perceived state of an intrinsic variable. An intrinsic variable is in the internal rather than external environment of the model; better, it is a variable that is not controlled by the perceptual control hierarchy. In real organisms these are probabaly (but not necessarily exclusively) physiological variables. The fact is, we don't know exactly what all these intrinsic variables might be; but, since starvation seems to be demonstrably a cause of reorganization in lab animals, one intrinsic variable is probably something like "blood sugar" or C2H34O67N12 (?). I don't know if there are any studies that directly monitored a physiological variable to see if rate of reorganization (however that was measured) is directly propotional to the level of the variable (or the difference between that level and some reference) -- but that is the kind of study that would be needed to flesh out the model.

>Bill claims that changing gain, if the gain is large enough, doesn't
>matter much. You say that is one of the aspects of learning.

It MIGHT be. All I meant in my post is that, given the PCT model, gain is one of the parameters of control that COULD be changed (when modeling reorganization). So could the perceptual function, comparator, etc. Maybe all at the same time, maybe not. All are possibilities within the architecture of PCT. But, as I said (or meant to), PCTers know very little about how real reorganization happens because there haven't been many studies of it -- and the existing literature on learning is a confusing mix of studies which involve no reorganization (like, operant conditioning studies) which are called studies of "learning" but are really just the exercise of existing control systems with changing feedback functions or disturbances, or some combination of reorganization and control.

>I ask, partly because it seems a very important issue in PCT

There are no "right" answers to your questions about reorganization right now because we have no standard "preparation" for studying reorganization. We don't know much about the phenomenon (not nearly as much as we know about control -- where we do have a standard "preparation" -- the tracking task).

> One aspect of the
>study is the learning algorithm. Do we change the perceptual weights
>so that the Little baby knows what patterns to try to perceive? That
>may be necessary, but I don't think it will be the whole story

I think Bill's arm model does incorporate a learning algorithm -- it tunes up the control system as it behaves. As far as what your learning model "should" do to reorganize -- that depends on what you are trying to get it to do. There is no PCT dogma on reorganization yet.

> Do we use Genetic Algorithms on ECSs working in imagination
>mode to organize patterns of reference signs that create cooperation and
>not much conflict?

Sounds good to me.

>   There are all sorts of possibilities.

You bet.

>   I had hoped
>that those who have worked with PCT for many years would by now have a
>good handle on what works and what doesn't in training the system.

Sorry. I've spent most of those years trying to show how control works. It's been virtually impossible, so far, to convince many people that behavior is control, let alone that "learning" is a process of developing new ways of controlling. But I think Bill has developed some useful "training" algorithms. I'm sure he'd be happy to explain them to you.

>And that word "training" brings up another aspect: teaching. What happens
>when a hierarchic control system is taught rather than learning from
>experience?

Great question. We can speculate but what we need is research!!!

>Is there a party line on the difference between learning from experience
>(acquisition), and being taught?

Nope.

>Being taught surely cannot involve
>random restructuring in the face of excessive intrinsic error, can it?

Not necessarily. Teachers certainly can give some direction (not always helpful -- that's why it would be nice to learn about it) to reorganization. PCT research on teaching is needed -- graduate students take note!

>Puzzlement.

You bet. I think that our best bet for finding some answers about re-organization is to ask nature -- not the PCT model. The questions we ask, however, could be based on PCT.

Regards

Rick

Date: Mon Feb 03, 1992 8:26 pm PST
Subject: BEERBUGS; emotions

[From Bill Powers (920203.1000)]

Pat & Greg Williams (920202.1043) --

> If the feet are not slipping on the ground, the legs will tend to move
>slower than if the feet are slipping, because the physical forces
>opposing the forward movement of the bug are greater in the former case
>(the bug will move faster relative to the ground). With feet slipping
>and therefore faster leg movement relative to the body, the backward
>angle sensors will be activated with a smaller cycle time -- and those
>little feet will fly. In Beer's model, he didn't really put in all of
>the F=ma physics.

It looks to me as though he didn't put in ANY physics. Where in the model are the "physical forces opposing the forward movement of the bug?" Is there anything but the neural driving signals that affects the speed of leg movement?

Are your predictions about the legs moving faster on a slippery surface based on intuitions about real bugs, or on what the model, as now set up, would actually do? I don't think you could even put a slippery surface into the simulation. Have I missed something?

I read recently (Science News, I think) that when threatened, cockroaches rise up onto their rear two legs and run forward at very high speeds -- I think it was up to 35 mph or kph -- anyway, FAST. If Beer's model is really based on the complete circuitry of the cockroach's nervous system, it should be able to do the same thing. Unless, perhaps, the circuitry isn't complete.

Somehow I feel that the circuitry is based more on connectivity traced through the main pathways, with puzzling connections (like recurrent collaterals and sensors that don't seem to do anything) just left out.

Can Beer's model even vary the speed of walking?

I guess I'd better buy Beer's book. Would you post the ordering info again?

RE: Open loop control

>Several invertebrates have "escape" (demonstrably feedforward) circuits
>which, in "reaction" to ominous sensory input, get the animals the hell
>away. ANY WHICH WAY away!

As long as it's "away"? What is the effect of this "feedforward" reaction on the ominous sensory input, from the standpoint of the reacting system? Is it to make the ominous input even closer or greater? To have no effect? Or to reduce it? I think you're playing devil's advocate here. Words like "escape" and "away" and "toward" define actions by their consequences. In the traditional view of behavior, these consequences are just outcomes, and if they repeat then the outputs of the nervous system must have repeated -- the S-R non-sequitur.

Organisms that "respond" to noxious stimuli in such a way as to leave the stimuli unchanged or to make them greater don't survive. Only the organisms that control noxious stimuli relative to a reference level of zero survive. Only the control systems survive. Not so?

The "fact" that you cite about "demonstrably feedforward" circuits are cited from the standpoint of the traditional interpretation, which ignores feedback effects and control. I stand by my claim: there are no open-loop "reactions."

>Local "feedforward" can be embedded in global feedback loops. But
>that doesn't make it local feedback!

That doesn't make it "feedforward," either. It's just an output that depends on the input to a function. The same is true of input-output relationships in perceptual functions and comparators -- that doesn't make them "feedforward" functions. I hate the term feedforward because it's just a verbal play on "feedback," and doesn't have any comparable special meaning or refer to any fundamental principles. It's another example of letting words push us around. Why not "feedsideways?" Why not (to cite another of Pribram's contributions to wisdom) "overarching dependencies?"

There are cases, like the vestibulo-ocular reflex, in which the notion of feedforward is legitimate -- the sensor operates the effector with no direct feedback effect. But I'd rather use a circumlocution than encourage the use of "feedforward," which is generally used as a way of sneaking S-R theory back into the picture and showing that feedback is

just a minor side-issue.

Kent McClelland (920202.1401)

>Are emotions simply epiphenomena of control or lack of
>it, or do you see them as playing a more significant role in the
>hierarchical model?

We know about emotions because we experience them. Anything we experience must be a perception. So emotions are perceptions. What is it that we perceive when we perceive an emotion?

It's pretty well established that among the things we perceive are somatic states: heart rate, respiration rate, adrenalin or its sensed effects, and in general shifts in the pattern of biochemical organ-system activity in the body. We can recognize different patterns in these sensed states ranging from excitement and intense preparation for action to that zingy sense of well being to a sort of normal neutral to apathy to depression and sickness. I would place such identifiable patterns at the configuration level of perception.

Emotions also involve perceptions of the outside world and our relationships to it, and cognitions. When we feel anger, we feel the heightened state of preparation to act, but we also feel that this emotion is directed toward things outside us: we're angry AT something or somebody. Fear involves an essentially identical bodily state, but the feeling seems more directed toward going AWAY from something, escaping from it. So the ordinary perceptual hierarchy is involved, with the somatic perceptions merging into our perceptions of other kinds.

The easiest emotions to understand are those that can be identified naturally with error signals. When you're angry, you want to take forcible action to remove a disturbance. When you're afraid, you want to take action to get away. So there's clearly a goal involved: attack, or flee. This is a *behavioral* goal, but it seems to entail feelings, too.

The obvious answer is to say that somewhere around the level of configuration control, which is located in the midbrain, I think, the hierarchy splits into two hierarchies: one behavioral, using the muscles, and one somatic, using the biochemical organ systems. The outputs of higher systems alter reference levels for both the action systems and the biochemical systems. The biochemical branch of the hierarchy works through the hypothalamus and thence through the pituitary, where reference signals for all the major hormone systems are set. There is also innervation going directly from the brainstem to every major organ, again presumably setting reference signals (practically any biophysiologist could do better at this level of guessing than I can). These reference signals specify changed states of the somatic systems, which we sense as the feeling-component of emotions.

So the general picture is this. In controlling higher levels of perception (either initiating actions or resisting disturbances), the brain uses motor systems to produce overt action, and at the same time adjusts the biochemical systems as appropriate to the action. The lower-level actions, behavioral and somatic, are driven by error signals in the brain's hierarchy. The resulting perceptions are partly feelings from the body corresponding to the biochemical controlled variables, and partly perceptions of the ordinary outside-world and kinesthetic kind.

Therefore all control actions involve emotions, because what we call emotions are made of the behavioral and somatic perceptions involved in normal behavior. All behavior has a feeling component.

We don't normally give emotion names to the feelings that accompany such actions as tying a shoelace or sharpening a pencil, although there is always some feeling state in existence. We reserve emotion-names for situations in which the feelings depart considerably from their normal or neutral state. And that means, mostly, situations in which there are large errors. Only then is the somatic component of action sufficiently unusual to be perceived as anything special.

One way to create abnormally large errors is to put two control systems into conflict. This greatly reduces the amount of motor action that will take place, and creates abnormally large error signals in at least two control systems. The error signals that would be converted into overt action are ineffective because of the conflict: the muscles don't actually come into use. We want to attack, but don't; we want to flee, but don't.

However, the error signals also enter the somatic branch of the hierarchy, and as usual reset reference signals in the biochemical systems, preparing appropriately for the actions. If both conflicted systems would call for strenuous physical action, both put the body into a heightened state of preparation. But nothing happens, because of the conflict. This is the state in which I think we are most likely to say we are feeling a strong emotion. We are physically prepared for the action, but the action doesn't happen, or doesn't happen normally.

The term "emotion" is pretty much of a grab-bag and I wouldn't venture to try to explain every situation in which that word might be used. But I think we can get pretty far by saying that every emotion is associated with an error signal, and therefore with a goal. Many emotions can be interpreted by asking not what it feels like but what you want to do when you feel it.

Of course the CT interpretation is not that you want to do something because you feel an emotion, but that the emotion arises in the course of wanting to do something. The want comes first. If the feeling arises because of a disturbance, the disturbance implies a goal, and the departure of something from the goal-state creates the error signals that give rise to both emotions and the desire to act. You are most likely to call the accompanying feeling an emotion if the action is blocked or ineffective, so that the error reaches an abnormally large magnitude.

The "good" emotions are harder to explain. One explanation is that we give "good" names to feeling changes that occur as large errors are reduced. Normally, when behavior is effective and we control our perceptions successfully, we don't feel emotions. We may feel serene or competent or good or whatever, but we don't give dramatic names to the feelings that are present. Even the "good" emotions, therefore, imply that something has been wrong and is being fixed.

All this, of course, is by way of extrapolating logically from the theory and a minimum of facts. But you asked.

Best to all,

Bill P.

MESSAGE FROM MARY:
[from Mary Powers]
[Gary Cziko - from way back]

What I meant about ancestors is that certainly these discrete individuals exist, but there is unbroken continuity in the germ plasm - back to the beginning. Otherwise it's reasonable to think that mice are spontaneously generated from heaps of rags. In a bacterial culture, which one is the ancestor?

[Kent McClelland]

Bill has had two checks from Grinnell dated a week apart - one for "programs", one for "software". Is this a mistake or is one for last semester and the other for this semester?

Mary

Date: Mon Feb 03, 1992 8:27 pm PST
Subject: Re: Language & perception; BEER bug

[Martin Taylor 920203 11:45]
(Bill Powers 920201.0900)

>
>The slipper chewed the puppy
>The ball ate the ice cream
>The ice cream ate the ball
>
>I believe that Harris would say these phrases are unacceptable because in
>each case the subject of the phrase is an operator that does not take an
>argument, and so can't be connected by a transitive verb to an object.
>That isn't an a priori reason, but a generalization derived from
>considering many comparisons of phrases by individuals. The explanation
>should be stated more empirically: it has been found that ...
>
>The *a priori* reason I'm thinking of is simply that the implied events,
>which we can imagine, never actually happen. We have never seen anything
>called a slipper, a ball, or ice cream doing the sort of thing that the
>verbs imply. It's hard to imagine how they could. It isn't language, but
>the experienced world that tells us these sentences are unacceptable.

To a certain extent, you are probably right, but there is more to it, which makes the purely linguistic approach right as well. There is a paper on the neuroprose archive by Finch and Chater (Edinburgh) that illustrates why. What F&C did was to look at the contextual environment of all the entities in three different databases, and use statistical clustering methods to group those that occurred in similar contexts. The entities in the three studies were letters in text, phonemes in transcribed speech, and words in a large database (Usenet postings).

The "context" consisted of the distribution of entities of the same kind that immediately preceded the candidate entity, or of those that immediately followed it, or that occurred two before or two after the candidate. So the context profiles for each entity consisted of four vectors, each containing the probability of encountering each of the other entities in the lexicon (i.e. all the letters, all the phonemes, or all the words -- actually the very infrequent words were not counted, so one should really say "most of the words").

When they did their clustering, they found that at the coarsest level, words separated into more or less the normal syntactic classes (verb, noun, adjective, etc.), letters and phonemes into vowels and consonants. At more refined levels, words split out with more of what we might call semantic aspects in common within a cluster--adverbs of time, nouns labelling professions, and so forth. Remember that NOTHING in the data gathering made any reference to anything other than the neighbouring words. Perceptions had nothing to do with it at that level of analysis.

So, Bruce would be perfectly correct to say that the LANGUAGE contains the structures that determine that "The slipper chewed the puppy" is anomalous.

But then we ask "why do these words cluster together, and separately from those words?" Surely at that level, it must be because slippers seldom chew, whereas they are chewed more often. The structure of the language depends on the likelihood of the percept. Also, the likelihood of my oerceiving something (more subtle than whether the slipper is chewing) probably depends on whether that perception corresponds to a "proper" linguistic structure (I'm a bit of a Whorfian on this).

We certainly act as if the words have more impact than the perceptions. Last time, I mentioned dangerous words like "money" (and I don't mind getting into an economics discussion, though I don't see how PCT applies when the control systems are in individuals). This time, I might use "asbestos" as an example. There are two completely different minerals that share the name "asbestos" because both provide readily formed fireproof materials that look quite similar. One of these minerals is a significant danger for lung cancer, but the other is not (see the editorial in the March 2, 1990 Science). Most of the "asbestos" used in North America is of the safe variety, and yet some $10 billion per year is spent on removing it from public (and private) buildings. This senseless expense is based entirely on the use of a word in two senses at the same time

But consider. If we used the two different words for the two forms of asbestos (chrysotile and crocodilite), would they not have much the same distributions, linguistically? Would they not therefore be closely associated? And would people not therefore come to have much the same feelings about both? I leave that question open.

==================

\>

\>> ... introspection is a suspect way of studying what goes on in one's
\>>head, ...

\>

\>I don't know how else you could find out what's going on in your head by
\>way of perceptions.

\>

It was the reliance on introspection by psychologists of the late 19th century, and its failure to provide any significant insight into mental processes, that led directly to Watsonian behaviourism. I don't trust it, even when it is my own mental processes I examine. How is it determined what aspects of the mental processes come to consciousness, and what does it mean that those aspects are consciously available? I can tell you what I consciously recognize to seem to be happening, but what is really doing the work is unknown to me. But for what it's worth, yes, I do need imagery to solve simultaneous equations or to make algebraic substitutions.

Introspection is OK, I think, if it directly contradicts some statement about mental function, such as "No one can image a scene they have not experienced." But much beyond that, it can give only vague clues as to what is going on.

Martin

Date: Tue Feb 04, 1992 3:23 am PST
Subject: BEER'S BUG, etc

[From Rick Marken (920203)]

Pat & Greg Williams (920202) say, in response to my complaint that Beer's way of constructing bugs is like doing arithmetic with Roman numerals:

>Nobody said it would be easy (or "cognitively penetrable"!), Rick. PCT's
>"functional-level" models appear easier to grasp, but all of the nitty-gritty
>details will have to be supplied to make them work in the real world.

Agreed.

> when your organism and environment
>are entirely conceptual (software), handling the low-level stuff is TOO EASY
>because it is solely representational (in the CP sense). But what can one
>build out of fairly realistically modelled neurons only? Beer's bug, for a
>start.

Agreed.

>We're working on an expanded Nervous System Construction Kit which will allow
>PCTers to apply their ideas to building nervous system models, and we hope
>they will do so.

All right, you guys win (sorry Gary, I tried my best). I think Beer is trying to do something a bit different than we (PCTers) are; he is trying to see what you get when you construct systems out of components which neurophysiologists say actually exist; PCTers are more interested in what variables systems control -- though we care about how neurophysiology might support or constrain that control. I think it would be VERY useful to have a NS Construction Kit that made it easier for PCTers to construct models out of "real" components. What might be nice (maybe this is what you are doing) is to have a FNS (functional nervous system) set which would let you make perceptual functions, output functions and reference inputs (to comparators). And then this FNS could be "compiled" into Beer's RNS (real nervous system) components. Easier said than done, I'm sure. But possibly of ENORMOUS value for modellers, researchers and teachers.

By the way, GREAT job, Greg (and Pat?) on Closed Loop; beautifully produced. I was also thrilled to hear that you are working on my book too. I am really glad to have you folks doing it; you do EXCELLENT work.

Best regards

Rick

Date: Tue Feb 04, 1992 3:26 am PST
Subject: Re: language learning

[Martin Taylor 920203 11:30]
(Joel Judd 920203 10:24)
>
> I'd like to repeat back the three
>types of change with examples to check my understanding. I'll use the old
>standby, phonology.
>   ------- I--------
>>1) a change in the characteristics of the sensory function that maps inputs
> into perceptual signal; >such reorganization changes the "meaning" of the
> perceptual signal.
>
>I assume this means that given the same environment, the only thing
>changing is how I percieve it. Is this as simple as old psych visuals like
>a Necker cube or the old lady/young woman?
>
I don't see it that way. The problem of ambiguous figure preception is
what I would call "Perceptual conflict." But Bill doesn't agree about
the existence of perceptual conflict, so there is a problem.

Regardless of that, I think that what we are talking about is only the
readjustment of the weights at the input to a simple "add-and-transform"
node of the kind that appears in most neural network studies. It's kind
of a redescription of a perceptual feature. At higher levels, the perceptual
function is more complex, and the changes may result in more drastic
alterations to the meaning of perceptual input ("meaning" being aligned
with the concept of what you do about the input). But the idea of
switching back and forth between permissible inputs is not part of this
aspect of learning. To do the switching, one must already have learned
the two possibilities. It's a different phenomenon.

I'm still not clear what is meant by "Intrinsic Error."

Martin

Date: Tue Feb 04, 1992 10:42 pm PST
Subject: Re: Reorganization

[Martin Taylor 920204 12:10]
(Bill Powers 920203.1900)
Bill, where do you get your time stamps? I don't think I was even awake
at 920203 0926. Is it the arrival time of the mail at your node? If so,
it makes identifying the relevant message a bit difficult in a later
backtrack. Wouldn't it be better to use the identifier time entered by
the author?

On "perceptual conflict": I think you are doing what you complain about--
using a verbal equivalence to assert a real-world equivalence. I take
"perception" in "perceptual conflict" to refer to the output of some
perceptual combination function in an ECS, and you take me to task because
the "conflicting" percepts are not simultaneously in consciousness.

I do think the Necker Cube and its analogues represent perceptual conflict in the sense of percept being some function that has been developed of the input. When the cube is switching between consciously perceived states, each of the possibilities has been fully developed, and there is a conflict as to which is allowed to be consciously perceived (I believe that this stems from a basic conflict as to which of a set of ambiguous percepts is to be allowed to affect behaviour, but that's an extra, not relevant to this particular discussion).

In 1962-3 I did some studies with a couple of students on the nature of reversing figures, using the timing relationships among reversals, and the relationships between the number of different forms seen and the number of switches reported. We were able to determine that the reversals were not a response artefact (we had biased the subjects to reports one kind of form and not another, or to report both kinds), and much more interestingly, we were able to find a figure that seems to have only two perceived states (almost everyhting you can draw has a multitude of percepts if you don't tell the subject that there are only two).

Using the figure with two states, we measured the distributions of times between reversals over long durations (four, 36-minute sessions during a one week period). We found that the distributions were quite precisely modelled by a random walk process over a set of S micro-states, each of which could be set to one percept or the other. If more than K states were in percept 1, then 1 would be consciously perceived. If less than L states were in percept 1 (L < K), then percept 2 would be reported. If between L and K states were in percept 1 the conscious percept would be whatever it was before it got into the "dead band". The micro-states changed randomly at a rate k.

This could be seen as using many parameters to fit an arbitrary function were it not for some fascinating facts. Firstly, there are four parameters that must fit the shapes of two survival curves. Second, over the experiment there were 16 such curve pairs per subject, and the number of states S remained the same for a given subject for all the fits, while the rate function k grew reasonably smoothly over all runs. Thirdly, for quite drastic changes in the survival curves that happened sometimes between 9-minute mini-runs, the fitting difference was always exactly one unit in K or L or on one occasion both K and L in the same direction. We came to the conclusion that the numbers S K and L really represented somthing in the subject's perceptual mechanism. One of our subjects had 32 individual perceiving functions devoted to the problem, the other had 35 (if I remember the numbers correctly).

If we identify each microstate with the input perceptual function of an ECS, the conscious perception would be derived from a hysteric majority function. I would call this "perceptual conflict resolution" to produce the conscious perception.

Martin

Date: Wed Feb 05, 1992 12:31 am PST

From Pat & Greg Williams (920204)

From Bill Powers (920203.1000)]

>It looks to me as though he didn't put in ANY physics. Where in the model
>are the "physical forces opposing the forward movement of the bug?" Is
>there anything but the neural driving signals that affects the speed of
>leg movement?

>Are your predictions about the legs moving faster on a slippery surface
>based on intuitions about real bugs, or on what the model, as now set up,
>would actually do? I don't think you could even put a slippery surface
>into the simulation. Have I missed something?

Beer's model of the physics is a proportionality constant between bug velocity (not acceleration) and leg force. Slipperiness would reduce that constant of proportionality. Beer didn't try that, but you could, using NSCK.

>I read recently (Science News, I think) that when threatened, cockroaches
>rise up onto their rear two legs and run forward at very high speeds -- I
>think it was up to 35 mph or kph -- anyway, FAST. If Beer's model is
>really based on the complete circuitry of the cockroach's nervous system,
>it should be able to do the same thing. Unless, perhaps, the circuitry
>isn't complete.

Beer didn't put in the escape circuitry (which typically is seen following puffs of air applied to the posterior region of a cockroach). He DID put it in in a later model (see AMERICAN SCIENTIST, Sept.-Oct. 1991). But still, his bug's nervous system model isn't "complete." How could it be with only a few dozen neurons? But that isn't the point. The point is to capture the IMPORTANT things in the model and to leave out the redundancies, etc. If something deemed important later was left out, then attempts are made to add it.

>Somehow I feel that the circuitry is based more on connectivity traced
>through the main pathways, with puzzling connections (like recurrent
>collaterals and sensors that don't seem to do anything) just left out.

Perhaps. But it is at least arguable that a lot less faith and implicit ignoration (of coordinates and a lot more besides!) is behind, say, Beer's modelling of the walking circuitry of the bug than is behind, say, your own models of high-order loops affecting lower ones by altering the latter's reference signals.

>Can Beer's model even vary the speed of walking?

Yes, "directly," by varying the output of an activation-level neuron which excites the gait circuitry. In NSCK, inject current in LCS cell in the L3R3 setup.

>I guess I'd better buy Beer's book. Would you post the ordering info
>again?

Academic Press has a toll-free order line, 1-800-321-5068. Beer's book has ISBN 0-12-084730-2.

>As long as it's "away"? What is the effect of this "feedforward" reaction
>on the ominous sensory input, from the standpoint of the reacting system?
>Is it to make the ominous input even closer or greater? To have no
>effect? Or to reduce it?

To negatively feedbackly reduce it, of course. The local feedforward is part of a larger feedback loop, as Rick argued way back.

>I think you're playing devil's advocate here.

Well, SOMEBODY has to keep interesting discussions going so there will be stuff to put in CLOSED LOOP....

Best of luck to Gary in his encounter with Dr. Beer. We hope nobody is driven buggy!!!

P.S. to Gary: You should write to Pocket Soft, Inc., makers of "RTLink" and "RTPatch" (the RT stands for "Run Time") software, and tell them they SHOULD be "RTlink" and "RTpatch"! Maybe you could turn this into a hobby....

Pat & Greg

Date: Wed Feb 05, 1992 4:07 am PST

[From Bill Powers (920204.1830)]

Martin Taylor (920203.1700) --

A very nice, informative, lucid post that gets us somewhere.

>... Bruce would be perfectly correct to say that the LANGUAGE contains
>the structures that determine that "The slipper chewed the puppy" is
>anomalous.

I agree that it does -- I'm not accusing anyone of making mistakes in their word-counts. The forms are there in language. But as you say, "The structure of the language depends on the likelihood of the percept." It would be astonishing if language suggested one structure relating meanings, while direct perception suggested another. That wouldn't be a feature of language, but a bug, given that language is for communicating.

I'm trying to get an idea into good enough shape to communicate it, and it seems that I can only handle one preparatory development at a time. Perhaps I should just try to say where I'm trying to go.

I think that the study of language *as language* can tell us something about perception, because in language there will be constructions that are there in order to communicate perceptions. I don't just mean objects and relationships, but higher-order perceptions, the ones that are hardest to pin down. But before we can use language analysis that way, we have to strip away apparent properties of language that are not strictly aspects of language at all -- for example, whether puppies chew slippers or slippers chew puppies. Basically, that kind of assymetry in language is trivial, and we shouldn't bother with it, except, as the cops say at the murder scene, for purposes of elimination.

But there is something important there, even with sentences concerning slippers and puppies. It's the way the ordering of the words tells us who did what to whom. Somehow "The slipper chewed the puppy" is telling us that the *agent* is the first noun, the event follows it, and the object comes last, in English coding. In HCT terms I would say that this sentence describes a relationship between one object and another object involved in an event, with the objects and the event explicitly named -- puppy, slipper, chewing -- and with *the form of the sentence itself* indicating the kind and direction of the relationship.

We can refer to this kind of relationship without even naming the event or the configurations and transitions involved in it: Subject, verb, object. That triad can be translated directly into a sense of relationship devoid of lower-order details. So it doesn't matter at all whether, when we substitute particular terms into the general form, the result is a likely sentence or refers to a likely percept. If we say "the quantum painted the ocean" we know what relationship is meant: it's the kind of relationship we call an operation, in which an agent does something to something else. You have to have two things and an action (the actions might be either a transition or an event, depending on whether it's ongoing or a temporal package). One thing is the agent that does the transition or the event that is received by or done to the other. The things are identified as we identify configurations, by some sort of description or name. The named transitions and events that constitute the verb involve the things in action, either as cause or effect (in this case). And the relationship involves the things, transitions, and events in a particular -- well, relationship -- with each other.

As Bruce has been saying, and I've been too far afield to acknowledge, the sentence form itself doesn't translate one-to-one into the relationship perception. But by putting the words specifically into a relationship, it can indicate, by convention, another relationship which is the one I am calling an operation. There isn't any word in the sentence that refers to the operation; the words refer only to the elements of the operation. But the sentence form as a whole does refer to the operation. There are conventional inversions of the sentence, such as the passive form, that allow you to emphasize the object by mentioning it first, but that form also points to the same relationship: object, passive verb (pointing the other way), subject. This lets you control two variables independently: designation of the relationship, and emphasis of one element of it. Other inversions and special terms let you express the same relationship as a question, with freedom to vary the emphasis.

By using the word "operation" here, I'm specifically acknowledging what may be important about operator grammar. It doesn't matter HOW Harris arrived at the idea of operators and arguments or at his classifications of various combinations. If these are indeed stable and discriminable forms, then they should have meaning to us even when we would judge particular examples to be grammatically wrong or otherwise unacceptable. We could figure out what the sentence is trying to say, although we might consider the meaning ridiculous or impossible or nonsensical. That doesn't matter. What matters is that the forms themselves are indicating some level of perception, perhaps relationships or perhaps other levels (language was not constructed with HCT in mind). We know what is meant by saying "The monkey blasted its mathematics" *at the level of perception that corresponds to the form*, even though we can't make sense of it at any lower level. The kind of relationship intended has been communicated.

But to make use of this idea that the forms themselves, stripped of lower-level meanings, have perceptual meaning, we have to divorce the forms from any lower-level content. Trying to bring lower-level meanings along with us as we analyze sentences for higher-level meanings simply confuses the picture -- it no longer matters what the lower-level meanings are. That's what I've been trying to get at in my roundabout way. Maybe I'm only now realizing what I've been trying to get at.

I think we have conventions for communicating a sense of category, a sense of sequence, a sense of program-like contingency. These conventions entail the use of special words and word forms, like "an" apple for categories, lists and terms like "then" and "next" for sequence, "if", "and", "or", "not," and "therefore" for aspects of programs. Again, I'm not suggesting any one-to-one correspondence; only conventionalized ways of referring to perceptions at different levels. We don't need the lower-level meanings to grasp the higher-level perception: if not argle implies bargle, and not bargle, then argle. Makes no sense at all, but the contingency is plain. The program level doesn't care what argle and bargle are: it just follows the rules. We can perceive the rule in the structure of language without even knowing what it's a rule about.

Now we get to the second part of your post: "We certainly act as if the words have more impact than the perceptions."

In the light of the above, we now have to be more careful about what we mean by saying "the words." What aspect of the words? "Asbestos," as the name of a configuration, means a fluffy substance we can use for insulation. We can still do that, even knowing that there are two kinds, and that using them has different implications. In statements using the word asbestos, such as "Asbestos causes pulmonary ailments," the meaning depends on whether the listener knows that there is an OK kind and a not-OK kind. The listener who knows this difference will hear the meaning at one level and understand it (the statement says that asbestos is related to pulmonary illness as cause is related to effect), and hear it at another level and judge it false (some asbestos causes illness; some does not; therefore the meaning of the statement that asbestos unqualifiedly causes illness is false).

You're right about the ways words lead us to actions. They purport, after all, to be descriptions of perceptions, and if the perceptions don't match what we intend or want to perceive, we naturally act to correct the difference. This, however, is not a property of language. It's the result of the way we convert from language to perception. The statement "Your house is burning" is linguistically impeccable, yet your house may not be burning. However, considering the risk of doing nothing if the sentence refers to a truly perceivable event, you may well cut your meeting short and go see if you can rescue anything or anyone. That's a logic-level decision. This is not a problem in linguistics -- you might do the same thing if you're daydreaming and get a sudden vivid image of your house afire.

The same consideration applies to the difference between chrysolite (sp correction?) and crocodilite. The words aren't important; they aren't what lead people to reject all asbestos including the harmless kind. What leads to the rejection is ignorance of a factual difference, and behaving to control the perception that the word "asbestos" means -- to the behaving person. The objective meaning doesn't come into it.

These are problems concerning the uses to which we put words, and how well we understand the process of using and interpreting words. Such things can be taught. What we're doing right now, and have been doing for months, is finding our way through the puzzles of words and meanings, so that we can create and use communication in a better way. From my standpoint, we're seeing more and more clearly that words establish only a sketchy and unreliable link from one person's meanings to those of another person (reliability declining as the complexity of the meanings rises). The meanings that the listener gleans from words, word relationships, word structures, and special forms, are always far richer and more detailed than the words justify, and can often suggest perceptions quite different from those intended, and quite at variance with accepted knowledge about the world. But that's a problem of living, not a special problem of language. We misinterpret many kinds of experiences that don't come to us through language, and are ignorant of many nonverbal facts, making the same kinds of mistakes that we do under the influence of language. When we learn not to make those mistakes in general, we will no longer make them using language, specifically.

RE introspection:

There are different kinds of experiences referred to by the word "introspection." One meaning, the one that prevailed while introspectionism was getting its reputation, is "imagination." Trying to find regularities in people's imaginations is a futile effort.

But another meaning is simply "noticing what is going on." You can say, introspectively, "I smell something burning," and unless you're deliberately lying this has to be a correct observation. The experience may or may not relate to anything going on outside you: perhaps you're about to have a seizure, or perhaps you're hallucinating or dreaming. That does not affect the correctness of the introspective report: you are experiencing the smell of something burning. Of course if you believe that your experiences invariably correlate with something outside you, then the implied assertion " ... and therefore something IS burning" is unwarranted. But if you are simply reporting appearances, whether they be located in the part of experience you call the "outside word" or in the other parts, you can't be mistaken, and your report is data.

I quite agree (in terms of my model) that we can't experience the neural functions involved in our perceptions. We experience only activities in the CNS, and even then, only some of them at one time, and perhaps some of them never. This isn't a problem for introspection; in fact introspection tells us the outcomes of these processes, so we can try to make models that could produce such outcomes.

On the other hand, we have to be aware that even our thoughts about matters like these are known to us only introspectively. While we can certainly know beyond doubt that we are thinking "some processes in the brain are not accessible to awareness," we cannot know beyond doubt whether this statement has any referent in reality. It can only be true or false in the context of a model.

I don't think that introspection gives us "vague clues." It gives us essential clues, if we are smart enough to figure out what they are telling us. I think that processes in our brains are more open to inspection than many have claimed. I think they are as open to inspection as the world we live in (and are, in fact, the same thing). After I understand what is open to inspection, I'll worry about what is hidden.

Best,

Bill P.

Date: Wed Feb 05, 1992 4:22 am PST

Subject: catching up on the mail

[From Wayne Hershberger--catching up on the mail]

(Martin Taylor 920120)
>In reference to Wayne's equation of 80 msec presaccadic
>psychophysical and physiological effects, I thought it relevant to
>quote much of the abstract from the thesis of a colleague of
>mine...Attentional Focus and Saccadic Suppression. R.G.Angus, York
>University (Toronto), October 1974.

Apparently you did not receive my previous post thanking you for your earlier mention of Angus's thesis. Thanks again, Martin, and thanks also for the extended quote. It helps confirm my earlier suspicions that Angus's research is only tangentally relevant to the issue Scott and I have been investigating. We are timing shifts in "retinal local signs" not elevated retinal detection thresholds.

****
(Martin Taylor 920131)
>We call it the "Little Baby" project, and I will describe it
>separately some time -- not now. The idea is to start with sensors
>and effectors much like Bill's Little Man (Arm Demo), and to use the
>concepts of distributed coordinated control that I posted a week or
>two ago to allow the Little Baby to go from a flailing infant to a
>controlled adult. I don't know if it will work. One aspect of the
>study is the learning algorithm.

Please keep us all posted on your progress. I, for one, believe that you are addressing the most important question that any control-theory modeler has to face; that is, what makes (keeps) the polarity of the feedback loops negative. The first thing to determine is the essence of negative feedback. I tried to start this thread with Bill Powers, and Rick Marken last year by asking them about E-coli (i.e., does the tumbling comprise a part of a negative feedback loop controlling the geocentric direction of locomotion, or does the tumbling randomly change the polarity of a control loop in which the feedback goes positive from time to time. Bill and Rick posted some very interesting replies and I should have pursued the matter further, but didn't. I am glad that you seem to be mining the same vein; it promises to provide a lot of gold.

****
(Tom Bourbon 920120)
>Can you give an estimate of how large? Large enough that it might
>produce a series of distinct evoked responses in human visual
>cortex? This June, I will be at the magnetoencephalography lab, in
>Galveston, Texas, where we might look for physiological evidence of
>the shift in humans.

Tom, the angular size of the phantom array depends upon the size of the saccade. The number of flashes in the array depends upon the frequency of the flashes. We've used frequencies up to 500 Hz and sacaddes as small as 5 degrees. The size of the array does not appear to depend upon saccadic masking (contrary to Martin's suspicions) because both the first and last flash painted DURING the saccade are visible). Your excellent idea of using magnetoencephalography to look for the neural substrate of the phantom array and its precursors is intriguing. Where would you look, in the frontal, parietal or occipital cortex? Scott Jordan may have an opportunity to do something along these lines. Scott has been awarded an Alexander von Humboldt award to work for one year in Hans Kornhuber's lab in Ulm Germany (a post doc). You may recall from your reading of Kornhuber's chapter in _Volitional_action:_Conation_and_control_that Kornhuber uses such exotic techniques. Some guys have all the fun!!

****
(Bill Powers 920120)
>My subjective experience is that form recognition moves with the
>locus of attention in the visual field and is limited in angular
>range. Things outside the circle of attention really aren't seen as
>configurations. This idea might be experimentally testable -- Wayne,
>what do you think?

Researchers investigating visual attention tend not mess with preattentive processing whereas those investigating preattentive processing tend not to mess with visual attention. I don't know either research literature well enough to answer your question. If someone has an answer, I'd be glad to listen because the question has implications for the selection/identification of saccadic targets, an interest of mine. For instance, is visual attention for the purpose of identifyting configurations as saccadic targets, or is visual attention itself a response to a visual target configuaration? Or, are both of these alternatives true? Configurations do appear to serve as sacadic targets. Try counting the number of letters in the rows below by saccading from one to the next; the first row is much less difficult.

asdfghjklzxcvbnmqwertyuioplkjhgfdsamnbvcxzpoiuytrewqqazw

llllllllllllllllllllllllllllllllllllllllllllllllllllllllllll

****
ED FORD:
Congratulations on your TV program. I look forward to seeing it in Durango next summer, if not before. I am glad to hear of your meeting with Kathie Kolbe, and that she intends to join the CSG.

****
(Pat and Greg Williams 920126)
>... not needing to "represent" the environment is a Big Deal in
>current AI, since it has proven so difficult to model
>representations.

There is another problem with representationalism; it is solipsistic.

****
(Bill Powers 920126)

>To the bug, the signals are the reality -- just as they are for us.

Oh? What happened to the boss?

****
(Greg & Pat Williams 920127)

>where there is feedback, there is analogizing.

Perhaps you could exand on this point a bit; I think you may be on to something.

****

(Mark Olson 920128)
>At first I thought I was biased against nonliving control systems.
>But now it seems that I find the problem with that nasty
>environment/organism dotted line. The line seems a little more
>"real" for the living than the nonliving.   When the temperature
>drop of the room causes (ontologically) the furnace to go on, I can
>explain it with physics and the line itself becomes a designation
>without real ontological status. That line seems "more real" with a
>living system.

To me, the organism-environment interface is every bit as artificial (nasty) as the mechanism-environment interface.

Warm regards, Wayne

Date: Wed Feb 05, 1992 4:42 am PST
Subject: Reorganization

[From Bill Powers (920203.1900)]

Joel Judd (920203.0926) and Martin Taylor (920203.0926) --

Joel says:
>I have usually seen "Intrinsic Error" used in relation to physiological
>needs of the system. Yet learning is often discussed regarding
>activities that don't seem to be "life or death."

and Martin says:
>I'm still not clear what is meant by "Intrinsic Error."

Hokay.

Intrinsic error, which drives the process of reorganization, is just error that the organism knows about without any previous experience -- in other words, you don't have to have (i.e., the body doesn't need) an adult hierarchy of control systems (or any hierarchy at all) to recognize that an error exists. This requirement follows from the necessity that a reorganizing system exist and work from the start, so that the hierarchy can be acquired as appropriate to the environment that happens to exist around the organism.

For the same reason, intrinsic perceptual signals must represent variables that aren't the outputs of learned perceptual functions. Intrinsic reference signals must also be built in for the same reasons. The reorganizing system must work without any knowledge of the world outside the organism. That knowledge does not cause its operation; knowledge results from its operation.

Now you know as much as I know. We can guess at some types of intrinsic variables. There are probably intrinsic perceptions related, for example, to hunger and thirst, although of course they aren't perceived as hunger and thirst -- they're just something important, and the reorganizing system acts when they deviate from the "right" state. I have proposed that error signals are also sensed as intrinsic variables (if large enough and chronic enough); this is legal because every control system, no matter what it controls, contains an error signal. Comparators, being elementary functions, can be built in, not requiring learning in order to exist. This means that losing control (and thus producing large error signals) or not having it in the first place, can consistute an intrinsic error condition, regardless of what is being controlled. The latter condition -- "regardless of what is being controlled" -- is the essential one if reorganization is to be effective from the beginning of life.

It's possible to do research aimed at identifying intrinsic variables. When an intrinsic variable deviates from its reference state, the immediate effect on behavior is not a tendency toward behavior that will correct the intrinsic error, but a change from smoothly organized behavior to random variations. That's the sign of reorganization. As reorganization is not systematic, the initial effect could well be to make the intrinsic error even larger. But that will speed up reorganization and the next change will come sooner, so temporary organizations that increase intrinsic error will not last as long as those that decrease it. In the long run there will be a biased random walk to the state of zero intrinsic error, if the organism doesn't die first. The last organization present at the time intrinsic error goes to zero is the one we say has been "learned." It is better to think of it as the one that prevents further reorganization. At this point, the proposed intrinsic variable would be at its reference state.

There is no logical connection between the control system that results from reorganization and the intrinsic error that causes reorganization. This is what makes reorganization so powerful. If it's necessary to learn how to align a pointer an inch to the left of a moving target in order to reduce the state of hunger to zero, that control system with that reference level will be acquired, if reorganization beats out starvation.

Skinner found that "shaping" gives reorganization a much better chance of working than simply withholding reward (maintaining intrinsic error) until the requisite pattern of behavior appears. If the change in organization required to restore control is small, the chances are much better that a random process will discover it than if the required change is drastic. So the experimenter allows a new behavior to reduce intrinsic error (produce reinforcement) when the new behavior represents a partial step toward the ultimate form that the experimenter wants to see. This greatly increases the chance that a blind variation will succeed at each step, so that very great changes can be brought about through a series of small steps.

When intrinsic error is zero, the rate of reorganization does not necessarily drop to zero. I have postulated, furthermore, that focusing attention on some part of the hierarchy directs to that part whatever amount of reorganization is going on. If the zero-error rate of reorganization is not zero, then wherever you focus your attention, some degree of reorganization will be occurring. If the systems involved are already organized for a minimum of intrinsic error, reorganization will just cause variations that keep returning to the same state. If the environment has changed its properties slightly (inducing error and thus attracting attention), the reorganizing will converge to a new form not far from the old one and then vary around that form if the zero-error rate of reinforcement is not zero.

A lot of what we call "learning" is not reorganization, but memorization. When I show you the order in which to press the buttons to open the door, all you have to do is remember what you saw and create that sequence-perception again. This doesn't require any new capabilities or any changes in existing control systems. Learning vocabulary is probably learning of this non-reorganizing kind. A person who is a "natural mimic" probably doesn't even have to reorganize to produce the sounds of a new language -- there aren't any sounds this person doesn't already know how to produce. All that has to be "learned" is which sounds to make, and that requires only memory.

There are probably learning algorithms too, which, once acquired through reorganization, simply have to be applied to work out the correct action in a new circumstance, like a computer program. The algorithm, of course, is a systematic process, and so is not reorganization.

Joel:
>... it has also been suggested and promoted by Gary that Blind Variation
>and Selective Retention (BVSR) process outlined by Campbell seems to be
>a promising description for the reorganization process

Yes, I've suggest it too. The "selective retention" part doesn't require a separate "retention" mechanism; an organization is "retained" when it is effective in correcting intrinsic error and stopping the process of reorganization.

RE: language

>What's left across people? Whatever aspects of perceptual control are
>MANIFESTED IN language.

Couldn't have said it better myself.

Martin Taylor (920203.0926) --

>The problem of ambiguous figure preception is what I would call
>"Perceptual conflict." But Bill doesn't agree about the existence of
>perceptual conflict, so there is a problem.

Perceptual conflict implies that both sides of the ambiguity are being perceived at once, leading to incompatible actions. But I don't think that "conflicting" perceptions like the Necker cube or the stairs or the faces and vases are actually perceived at the same time. If there's any conflict, it would result from trying to perceive both aspects of the input at once. A parallel example: look at the back of your left hand. Now look at the palm of your left hand. Now look at them both.

You can't see both because the action you take to see one rules out seeing the other. In the Necker cube, you can see the cube only one way at a time: the "action" required to flip to one makes seeing the other impossible. I think it likely that this indicates a cross-connection between the two perceptual computations, making them mutually exclusive. Maybe the same machinery is involved, but with a parameter change that create either the one output or the other but never both.

"Perceptual conflict" is really a metaphor, isn't it? Perceptions are just reports; they don't try to do anything. If you add two perceptions together, you just get the sum, a new perception; If you OR them together you get a superposition, like a double exposure. If you wanted the sum, or the double exposure, you wouldn't have any conflict about it. And anyway, it wouldn't be the perceptions having the conflict in any case; the conflict would be between the systems receiving these perceptions and trying to do incompatible things with them. I think that "perceptual conflict" is the same sort of thing as "stimulus generalization" -- it attributes the result to an agency of the wrong kind. A stimulus can't generalize, although an organism can. Perceptions can't be in conflict, although the control systems controlling them can.

From an earlier post on adaptive models:

> One aspect of the study is the learning algorithm. Do we change the
>perceptual weights so that the Little baby knows what patterns to try to
>perceive? That may be necessary, but I don't think it will be the whole
>story...

A self-reorganizing model has to have built-in criteria that tell it when it needs to reorganize some more. These can be any criteria that are affected by the behavior of the Little Baby. A true reorganizing system wouldn't care what the Little Baby ends up perceiving and controlling, as long as perceiving/controlling in a certain way corrects intrinsic error. If you want the Little Baby to perceive/control in some predetermined way, you have to make intrinsic error depend on departures from that way. I think it would be interesting NOT to try to force a particular mode of perception and control, but see what the model finds for itself.

My adaptive arm model actually uses error signals in the arm control systems, both proportional and rate, to stimulate the changes that optimize the arm. This is NOT part of Version 2 (with dynamics) but will probably be in Version 3. I've had it running. Actually my method is not reorganization; it's systematic, and I propose it as the way the cerebellum stabilizes the limb control systems. I haven't tried a real reorganizing system yet, but will, perhaps as an option for Version 3.

> Do we use Genetic Algorithms on ECSs working in imagination
>mode to organize patterns of reference signs that create cooperation and
>not much conflict?

If you use them in imagination mode, you lose the advantage of having real environmental properties in the loop (unless the imagination mode perfectly simulates them). I think it's better to have reorganization working all the time, on line. The critical thing is what conditions of the system are monitored by the reorganizing system as intrinsic variables (independent of learning).

Best to all,

Bill P.

Date: Wed Feb 05, 1992 10:03 pm PST
Subject: social bonding

Certain aspects of "male bonding" involve pretended threats of violence. I understand that this is common to all mammals. Bateson talks about it in an essay on the origins of play reprinted in _Steps_.

How can a mammal without language communicate something about its relationship with another? How can it communicate "I am friendly"? Basically, by initiating what appears to be a threat and then not carrying it out. Watch puppies playing. They growl and grab with their mouths, but don't really bite. This is what we call play.

We humans continue in the same mammalian boat. (Perhaps it's a bigger ark than just mammals, but that doesn't matter just now.) Why? We do have language, why do we need these primitive, pre-linguistic forms of communication? Because we can lie with language. But it is much more difficult to lie with what we call body language. And people who learn to manipulate their image, their presentation of self, evoke profound distrust when we find them out--salesmen, politicians, even actors. But we do find them out because of inconsistencies. Deception in nonverbal communication is very hard to sustain. We appear to control for information about our relationships with others. When there is incongruity between consciously controlled verbal channels and other channels of communication, a controlled perception of some sort is disturbed.

Deception is hard to sustain because when nonverbal communication occurs naturally we control it unconsciously. When we attempt to carry out consciously control that is usually unconscious, we interfere with it. We stumble, or the bicycle falls over. We can rehearse scenarios (like President Shrub), but circumstances don't always follow the script. Hence closely managed "photo opportunities."

The feigned aggression of play can easily escalate into real conflict. There are many accounts of people becoming great buddies after having an introductory fight. We might each be able to share some--men and women alike. The two people didn't like each other at first. Why? I think because they had so much in common. Perhaps they each saw the other as a competitor for the same interests.

There are other ways to communicate relationship. Deborah Tanner's work suggests that a way prevalent among women is to sacrifice something of value for the sake of the other. This is called accomodation, or compromise. Person A wants arrangement (a), and person B wants arrangement (b). A ascertains what B wants, makes sure B knows that she, A, wants (a), or at least that she wants something different from (b), then changes her goals so as to include (b) instead of (a), and makes sure that B knows that she has done so. B is then appreciative, and reciprocates when the opportunity arises, or else A expresses hurt feelings, anger, etc.

This works OK when (a) and (b) are token sacrifices, and the relationship itself is the real subject of nonverbal discourse. The process degenerates to manipulation when the expectation of B reciprocating some specific thing that A wants becomes the motivation for A's sacrifice.

There was an "I cut, you choose" strategy proposed for arms reduction, before Gorby performed his astonishing disappearing act. It's worth keeping around for other purposes. Each side puts things to be exchanged on the table. Each side assigns values to all the things, from both sides. The values are private. They almost certainly don't match. A offers to give some (a) of relatively low value to A. B proposes several things that have the same or smaller value in B's value system. From these, A picks one (b) that has a higher value than A's (a) has in A's value system. They exchange. Each gives up something of less value in exchange for something of more value. The apparent paradox is due to the fact that their private assignments of value differ. (Detail: If A's proposed (a) isn't worth anything to B, or if A doesn't like anything that B offers in response, then A tries something else. They take turns initiating after each successful exchange. Other arrangements for turn-taking are possible, such as simultaneous offers.)

This dynamic often underlies successful use of the "female bonding" communication pattern Tanner describes. When it does, neither side sacrifices anything except in exchange for something better. Better in terms of her own value system, that is. It is remarkable and intriguing that each side can then be quite literally taking advantage of the other, with no detriment to either.

Sounds good to me!

Bruce Nevin
bn@bbn.com

Date: Wed Feb 05, 1992 10:22 pm PST
Subject: I got it Dag

[From Wayne Hershberger]

Dag Forssell:

I received the manuscript you sent me by conventional mail.
Thanks. The graphics are great.

Warm regards, Wayne

Date: Thu Feb 06, 1992 12:22 am PST
Subject: Unnecessary messages

The amount of traffic on this mailing list is oppressive enough. Fortunately I run my own system so that I have set up a separate account just to handle this bandwidth, but it still seems excessive.

Would it be appropriate to ask that at least personal mail like the following not be circulated to the entire list? A lot of the mail falls along a spectrum from general interest to very specific, and when only one or two individuals are likely to be interested in the mail, why load down the net with copies to everyone?

Also please note that some addresses, like the following, cannot be replied to. I suggest that users try to incorporate valid Internet or BITNET addresses in ther signatures if they can't get them into the headers.

>Date: Wed, 5 Feb 1992 08:38:00 CST

>From: TJ0WAH1@NIU
>Subject: I got it Dag
>Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.bio.ns.ca>
>To: Multiple recipients of list CSG-L <CSG-L@UIUCVMD.bio.ns.ca>
>Reply-to: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.bio.ns.ca>
>Message-id: <A4D0223A1685028D8A@AC.DAL.CA>
>
>[From Wayne Hershberger]
>
>Dag Forssell:
>
>I received the manuscript you sent me by conventional mail.
>Thanks. The graphics are great.
>
>Warm regards, Wayne
--
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: Thu Feb 06, 1992 12:40 am PST
Subject: learning

[from Joel Judd]

Bill (920204)

Bless you for the timely restatement concerning reorganization. It helped to refocus my attention on several aspects. For one, you said

>Tha last organization present at the time intrinsic error goes to zero is the last one we say has >been "learned"...There is no logical connection between teh control system that results from >reorganization and the intrinsic error that causes reorganization.

I interpret this to mean that the less constrained the blind variation, the more variety in learning "outcome." So in comparing 1st and 2nd language acquisition, one sees convergence in children because their learning coincides with CSH devlopment, and it is often focussed on particular linguistic aspects with which they have had prior non-linguistic experience (eg. reference). Older L2 learners, on the other hand, are coming at language acquisition with a lot of "baggage": a developed CSH and accompanying linguistic abilities. Is it surprising then that there is evidence for greater divergence among older language learners? The pre-linguistic experience and socialization drive that have shaped primary language acquisition are often gone and buried in the adult's reliance on language to define himself. The need to develop native-like skills in an L2 can be hard to come by. So what if I don't know the current slang, or the phonological contrasts, or verb inflections? I can rent an apartment, and go shopping, and talk to the natives.

I better quit before I talk myself out of a job.

>A lot of what we call "learning" is not reorganization, but memorization.

This is what struck me in reading Campbell's 1974 Evol Epist paper. For blind variation to be as pervasive as he proposes, there must be A LOT of primary learning going on, really fundamental perceptual stuff, early on, upon which social and envirnmental interaction can build. There must be a tremendous amount of change in the first 3-4 years. It just seems that so much of what we do as adults, the day to day stuff, isn't reorganization, but it's origins have to be. One example might be the following article

"The *decline* of visually guided reaching during infancy" by Emily Bushnell in _Bridges Between Psychology and Linguistics_, 1991. The message of this paper is that neonates (3-5 months) seem to employ ballistic arm movements in reaching-- movements that tend to be quick and inaccurate. Around four months, however, the movements become somewhat slower (.5 sec by one account) and "devious," or visually *guided.* Part of the discussion of implications relates to how learned reaching frees up attention for other cognitve activities and further development. I'd like others' reactions this article (a footnote says that the same paper can be had in _Infant Behavior and Development_, _8_ (1985) pp.139-55.

(If you check out the book, you get added bonuses like the first two sentences from the second paper: "A central aim of the behavioral sciences is to understand how behavioral sequences are centrally controlled. When we talk, type, or engage in other voluntary activities, we rely on a complex set of internal events to activate the appropriate muscles at the appropriate times.")

Date: Thu Feb 06, 1992 12:58 am PST
Subject: Reversals; BEERbug; word recognition

[From Bill Powers (920205.0800)]

Martin Taylor (920204.1210) --

Yes, I was using the time received. Will use posted time when given...

You will be glad to know that I have used a teakettle to blister two typing fingers. This should make my remarks significantly more pithy.

>I take "perception" in "perceptual conflict" to refer to the output of
>some perceptual combination function in an ECS, and you take me to task
>because the "conflicting" percepts are not simultaneously in
>consciousness.

You're proposing a perceptual model in which all possible perceptions (whether apprehended consciously or not) based on a given input are "fully developed." This would be the literal pandemonium model. OK. If that were the case, I should think that at least sometimes, under some conditions, one could attend to both perceptual signals at the same time and experience both orientations or interpretations at once.

My impression, however, is that it's literally impossible to see the Necker cube in both orientations at once. I think that the reason has to do with the fact that the "cube" is a two-dimensional object and we're trying to interpret it in three dimensions. In order to create the sense of depth we have to imagine a missing piece of information: the positions of the corners in depth. If we imagine that one corner is far away, we see one orientation. Changing the imagined depth information causes the cube to be seen oriented the other way. One signal has only one value at a time, so we can't imagine the depth dimension of a given point as being both near and far. (I just thought of that this morning).

By this hypothesis, we don't have two fully developed perceptions, but only one at a time. This would explain why we can't experience both at the same time -- the signals don't exist at the same time at the 3-D level.

This idea sounds consistent with your most ingenious experiment with reversals. As you say, most figures are susceptible to more than two interpretations. Each dimension of possible reversal corresponds to a piece of missing information that must be imagined in order to turn the 2-D figure into a 3-D figure. If there are K equally valid imagined (binary) states of the missing perceptions, then there should be 2^K possible reversal combinations, although a given person might not discover them all. For example, one easily-missed interpretation of the Necker cube is that of a flat non-regular hexagon. That's actually quite hard to see. The possible states may not be binary.

I was very interested in the rate-of-reversal results. This "random walk" behavior sounds a lot like reorganization, which as you know I associate with the focus of attention. It's as though the reorganization process were flipping randomly through all the states of imagined information that would make sense of the 2-D figure as a 3-D figure. If there were one "best" interpretation, this method would find it. This has got to have something to do with how perception gets organized in the first place.

Through a CSG member, Sam Randlett, I became acquainted with The Great Randy, a magician. He showed me a huge collection of "effects" he had worked out that depended on this reversal phenomenon, some with dozens of possible points of reversal with sensible interpretations for all of them. But one astonishing demonstration consisted of a large card crossed by many diagonal lines. By diverging or converging the eyes a little, the lines could be made to coincide in different ways, and one "best" way quickly stood out: it turned into a 3-D grid of rods!

Another demo was simple and eerie. Take a 3 x 5 white filing card and fold it in half so the crease is along the long dimension. Adjust it so the angle at the crease is about 90 degrees. Then stand the card up on a uniform table surface with the crease away from you, stand back a way, and look at it with one eye. With a little effort, you can make the card "lie down," so it appears to be oriented with the crease parallel to the table surface and the parallel edges resting on the table. Once you have made the card lie down, sway slowly from side to side, and prepare to jump out of your skin.

Pat & Greg Williams (920204) --

>Beer's model of the physics is a proportionality constant between bug
>velocity (not acceleration) and leg force. Slipperiness would reduce
>that constant of proportionality.

So v(body) = k * f(legs)? Accepting the linear force-velocity function, we still have a problem with those little legs. If the constant of proportionality is reduced, v(body) will become smaller, won't it? If the legs are fixed to the surface when in contact with it (as I deduce they are), the legs will move more slowly, not more rapidly, relative to the body. To make a slippery surface, you would have to allow the feet to move backward without moving the body forward. And then, because the output is a force, you'd have to add the way leg position depends on neural output when resistance to leg movement is reduced. How will the legs move when the bug is on its back?

From your description it seems that leg velocity is made proportional to the driving neural signal amplitude, and that body velocity is equal and opposite to leg (relative) velocity. Maybe we could bypass all this guessing if you would just post the calculation that relates body velocity to leg force.

>But it is at least arguable that a lot less faith and implicit
>ignoration (of coordinates and a lot more besides!) is behind, say,
>Beer's modelling of the walking circuitry of the bug than is behind,
>say, your own models of high-order loops affecting lower ones by
>altering the latter's reference signals.

Ooh, ow! Just remember that the neural circuits Beer omitted are those that didn't seem important from a non-CT point of view. If we can catch Beer's model in a wrong prediction, then it doesn't matter if the circuits he presented correspond to (selected) neural circuits: he's got it wrong. So far it looks as if his model predicts incorrectly when we introduce real physics into it. Unless cockroaches behave like his model would. It would be nice if Beer had based his model at least in part on behavioral data under appropriate external disturbances.

The reason I'm not hot about neural modeling is that it's just analog computing done with elements that purport to represent neurons. Calling the elements neurons doesn't change the computations. The connectivity of the real nervous system is certainly suggestive of the system's gross organization, and the idea of neural oscillators is useful, but the actual representations of neurons and the functions they compute is very limited.

I think we could have come up with the concept of neural oscillators without knowing the circuitry, and that we would have guessed that sensing positions of the legs would be important in creating finite back-and-forth movements. Even just guessing, we probably could come up with a walking bug that works quite like Beer's model. If Beer had an exhaustive model of all the neural connections, all the sensors, and all the functions computed by each neuron, he would have a model built from first principles. But I don't think that's what he has.

It looks as though you're forcing me to try my hand at your program. Sigh.

Bruce Nevin (Tue 92014 12:11:05) --

I hope my 920204 post lifts some of the fog we're plowing through. I know we have many basic agreements by now. I know also that there are weak spots in my use of the pandemonium principle, but so far I've seen nothing else that works better.

Here's one point we may be able to work through:

> ... I understand that an ECS for a relationship among words or a word
>sequence does not ordinarily accept non-word perceptions as input ...

Let's unpack this a little. Suppose there are some words recognized below the relationship level. They are recognized by specialized word-recognizers. On that we agree.

The outputs of these words recognizers are signals, not words. The signals indicate THAT the words are present in lower-level perception. They are not themselves words. You have to include lower levels of perception in your awareness to see the details that make one word-signal refer to something different from another word-signal.

Now suppose that these words are nouns referring to the experience of a red color, and of a square shape. The words are "red" and "square".

The experience of the red color and the experience of the square shape may also be present as signals. The signals are neither square nor red, but indicate that the experiences of squareness and of redness are occurring, if they are.

We now have four signals indicating the presence of squareness, redness, the word "square", and the word "red." These signals are all alike except for their sources.

Suppose these signals now enter a higher-level perceptual function. This function doesn't know where the signals came from. It has become organized to treat the signals in pairs: squareness OR "square", and redness OR "red" (I'm identifying the signals by their sources). This perceptual function will thus respond in the same way when presented with redness and "square", redness and squareness, "red" and squareness, or "red" and "square." The output stands for some relationship between redness and squareness -- for example, redness is perceived as an attribute of the squareness. The relationship "is an attribute of" is reported as existing. The relationship signal does not say what is an attribute of what; it merely says that attributeness is present.

The perceptual signal coming out of the higher system's input function simply indicates that the particular relationship is present. The same relationship might be derived from many other inputs to the same input function: it is totally redundant with respect to the sources of the signals, as long as the requisite relationship is present. And the output of the perceptual function contains no trace of the source of the signals.

There are now four ways to produce the sense of "is an attribute of". You can present a patch of red and the word "square," or any of the other three combinations. Or you can present any other words and experiences that are input to the same higher system and exemplify "is an attribute of."

I apologize for this hastily-constructed and awkward example. I hope that the principle comes through without objections to the details getting in the way. The principle is that higher perceptual functions don't care whether their input signals originated in a word-perceiver or in a perceiver for the nonverbal experience that occurs under the same circumstances that the word occurs. This way of viewing matters means that once we pass the level where words are recognized as such, there is no difference between verbal and nonverbal perceptions, and there is no specialization into word-handling functions and nonword-handling functions. If a sense of relationship is evoked, it does not contain information about the kinds of things that are related. That information exists only in the lower-level systems.

I'm ready to entertain any exceptions or modifications of this basic picture, but I think we have here a very simple substitute for the concept of "association." We haven't really added anything to the basic model of the perceptual hierarchy. We've just said that some perceptual signals are treated as equivalent to others if they occur under similar enough circumstances. We can now understand sentences like "U go B 4 me," or those children's puzzles where a picture of something substitutes for some of the words in a "sentence." We can even understand "purely linguistic" discourse -- all the inputs come from word-recognizers.

I'll stop with that -- what do you think? Is this getting somewhere?

Re: operator grammar

>No, Harris would not predict the unacceptability of these sentences from
>the operator-grammar description of them.

I didn't mean to imply that. I meant that Harris would say that the sentences (like "the slipper chewed the puppy") are rejected because they entail using a word like "slipper" as an operator taking an argument, contrary to findings about normal usage of that word. I didn't mean that he could predict normal usage.

I accept what you say about Harris, but let's leave that for another time.

I echo your kudos for the excellent job Greg Williams did with Open Loop. He weaves together many separate posts so the stitches are invisible. Really remarkable.

Best to all,

Bill P.

Date: Thu Feb 06, 1992 1:56 am PST
Subject: Re: Unnecessary messages

Bill Silvert (920205) says:

>The amount of traffic on this mailing list is oppressive enough.
>Fortunately I run my own system so that I have set up a separate account
>just to handle this bandwidth, but it still seems excessive.
>
>Would it be appropriate to ask that at least personal mail like the
>following not be circulated to the entire list?

While I probably wouldn't use the word "oppressive" to describe the amount of mail on CSGnet (I find much of it quite liberating instead), it is a LOT and so I feel that Bill S.'s request is appropriate.

Although personal messages are actually quite rare on the net, when they do show up it's probably because the sender doesn't have the recipient's address. So let me remind netters how to get addresses.

Just send the following command as the first line of a message to
LISTSERV@VMD.CSO.UIUC.EDU or LISTSERV@UIUCVMD.Bitnet.

rev csg-l

and you will get a list of all CSGnet subscribers and their e-mail addresses. Address with more than one part (separated by periods) after the @ are Internet address while those with only part after @ are Bitnet address and may require the addition of ".bitnet" to be reached from an Internet site.

BUT MAKE DOUBLY (even TRIPLY?) SURE THAT THIS COMMAND IS SENT TO LISTSERV AND NOT CSG-L OR YOU WILL DO EXACTLY WHAT YOU ARE TRYING NOT TO DO, I.E., SENDING "PERSONAL" MESSAGES OVER THE NET.

>Also please note that some addresses, like the following, cannot be
>replied to. I suggest that users try to incorporate valid Internet or
>BITNET addresses in ther signatures if they can't get them into the
>headers.

I believe that TJ0WAH1@niu.bitnet is indeed a valid Bitnet address. My software automatically adds .Bitnet when responding to Bitnet addresses and so Bill S. may want to consider doing this as well.

Finally, short personal comments don't cause nearly the same disturbance (at least not for me) if they are stuck into a post of general interest to the net. I'm not saying this is better than a direct personal message, but it is certainly better than a separate personal message sent over CSGnet.--Gary

==

Gary A. Cziko	Telephone: (217) 333-4382
University of Illinois	FAX: (217) 244-0538
Educational Psychology	Internet: g-cziko@uiuc.edu (1st choice)
210 Education	Bitnet: cziko@uiucvmd (2nd choice)
1310 South 6th Street	N9MJZ
Champaign, Illinois 61820-6990	
USA	

==

Date: Thu Feb 06, 1992 2:13 am PST
Subject: chewing the slipper

[From: Bruce Nevin (Wed 92015 13:03:56)]

Real quick reply (sorry).

(Bill Powers (920205.0800)) --

>This way of viewing matters means
>that once we pass the level where words are recognized as such, there is
>no difference between verbal and nonverbal perceptions, and there is no
>specialization into word-handling functions and nonword-handling
>functions. If a sense of relationship is evoked, it does not contain

>information about the kinds of things that are related. That information
>exists only in the lower-level systems.

Are you saying that there are no program perceptions for syntax, for the conventions by which words are arranged in linear order? Or are you saying that those same programs apply equally well to non-word perceptions? That Achumawi speakers arrange the nonverbal perceptions |chew|, |puppy|, and |slipper| in that order, but English speakers arrange them so that the perception |chew| comes after the perception |puppy|? The convention for linearizing operator-argument dependencies may differ from one language to another. In VSO languages, the operator word comes first. In an SVO language like English, the operator word "chew" comes after its first argument word and before its second. (Conventions of both languages provide for a few alternative secondary linearizations for emphasis of the word in front position, e.g. "*John* I can believe she dated".) Your proposal makes it difficult to see how this could be under program-level control. Similar problems with control of other aspects of linguistic structure above the word level. You can't very well say that the structure inheres in the world of nonverbal perceptions, to the extent that it is not universal.

Briefly, it appears to me that some higher-level control applies to words but not to their "referents" (nonverbal perceptions). Can you see a way past this?

>I meant that Harris would say that the
>sentences (like "the slipper chewed the puppy") are rejected because they
>entail using a word like "slipper" as an operator taking an argument,
>contrary to findings about normal usage of that word.

"Slipper" is not an operator. The difference between the two sentences is that the arguments of the operator "chew" are in the reverse order:

```
             N      Onn         N
        The puppy   chewed the slipper.
        The slipper chewed the puppy.
```

The two sentences are of the same form (same sequence of word classes). If we construct two sets of sentences with the same word choices, one in one order and one in the other, and then compare the relative acceptabilities (within each set) of corresponding sentences, we find that the acceptability-differences are not preserved under this proposed operation of reversing the order of the two Ns. Therefore, it is not an operation in the grammar: it is not a product of reductions and/or the assertion of additional operators.

We might suppose that "puppy" belongs to a class of animate nouns to which "slipper" does not belong, and "chew" belongs to a class of operators that requires an animate noun. Many efforts have been made along these lines, constructing systems of semantic features and rules for their combination. One essential problem with these approaches is the unnoticed fact that the semantic features are vocabulary items in a special metalanguage, of which the rules for combination constitute the grammar. Now: what is the semantic interpretation of this special metalanguage?

Perceptions, we could say, nonverbal perceptions in the control hierarchy. But why not just say that for the language, without recourse to the separate metalanguage? The usual answer is that the metalanguage of semantic features is regular and explicit (or would be if the program were ever successful, which it has never been so far), with a 1-1 correspondence of form with meaning, whereas the ordinary language is cluttered with ambiguity and irregularity. But those are precisely the claims made, with empirical demonstration, for Harris's unreduced operator-argument dependencies.

In this perspective, objection to the artificiality of some of his sources and intermediate stages seems kind of silly. There is no objection to the metalanguage of semantic features being different from ordinary language. It is the operator-argument dependencies we are after, as they are prior to linearization. Who expects them to "sound natural"? How could they?

The difference is that for Harris's semantic representation (unreduced word dependencies) there is in every case a step-by-step transition to ordinary language, using for each step only operations that are demonstrated to apply widely throughout the language, for relations between perfectly ordinary sentence-pairs as well as for relations of "daggered" forms to more normal forms.

That turned out to be less brief than I planned.

Bruce
bn@bbn.com

Date: Thu Feb 06, 1992 2:32 am PST
Subject: Re: Reversals; BEERbug; word recognition

[Martin Taylor 920205 13:45]
(Bill Powers 920205.0800)
>
>
>My impression, however, is that it's literally impossible to see the
>Necker cube in both orientations at once.

Here's the problem with introspection. I disagree, based purely on my subjective experience. So where does that leave us? There are all sorts of possible arguments, based on different theoretical positions. You may say that I am lying or deceiving myself, based on your own experience buttressed by your theoretical position. I say that I am not, because I perceive (consciously) what I perceive, and I know that I am telling the truth about it (but is that the truth?).

There is lots of research in different areas, indicating that it is the bringing to consciousness that causes the uniqueness of perception. I cited some of it at various times. Perhaps it is the preparation for action based on perception that both brings it to consciousness and that forces a choice among the possibilities.

You say that the flat hexagon is rather hard to see. But for a naive observer who hasn't been told what to see, it is not an unusal first perception of the Necker Cube. That was part of the point of the study--you perceive what you have been primed to perceive, at least consciously, and probably to a lesser extent unconsciously as well.

By the way, one form of asbestos is chrysotile, not chrysolite. I checked again with the Science editorial.

In your response to Bruce, you describe a situation with the words "red" and "square" and objects that have those attributes, and you assert that it does not matter whether the object attribute or the word results in a higher-level signal. Here is another case in which perceptual conflict can occur. It is called the Stroop effect. Imagine the word RED written in blue ink or in red ink. If you are asked to respond to the word as a label, by pushing a red button or a blue button, you will be quicker, and probably more accurate, if the word and ink colour agree. If they don't you get a conflict, resulting in slow and probably inaccurate responses. There are lots of theories about it. Timing relations between the colour presentation and the word presentation are important. The situation is even more complex when a higher-level abstraction is the source of a conflict. Imagine having to press a button marked "animal" or one marked "tool" and being presented with a cartoon horse with the word "hammer" written across it, and being asked to press the button according to which picture was being shown. You get one pattern of interference (the pictures are faster, if I remember correctly). But if the buttons are for Horse and Hammer, then the words are faster. Categories take more time to determine than their exemplars do, from the words, but the reverse is true from the pictures (I may have that reversed, but the principle holds, and I think I stated it the right way round).

There must be a place where the two sources of information converge to produce the same result, if you have to make an action choice based on one or the other, but they are equally clearly separate in their effects in many ways. Especially if one of the sources is word-based.

Martin

Date: Thu Feb 06, 1992 3:32 am PST
Subject: words and other perceptions

[From: Bruce Nevin (Tue 92014 12:11:05)]

Bill Powers (920130.0900) --

>no separate linguistic chain. It is all one hierarchy. Words are
>perceptions like any other perceptions: there is absolutely nothing
>special about them. Higher level perceptions are functions of lower level
>perceptions regardless of [what] those lower perceptions consist of
>. . . I recognize only one hierarchy
>of perception and control, and it must do everything that is done to and
>with perceptions.

I meant to specify, not separate control hierarchies, but lines of control or signal-passing within the one hierarchy.

I understand that the recognizer for the word "pattern" accepts only certain transition and configuration signals as input matching its reference signal. It is specialized for the word "pattern". It does not respond to the sound of a car starting, for example. In all general respects it is just an ECS like any other ECS, but in terms of its particular I/O connections and signals it is specialized for language, for English, and for the word "pattern." Similarly, the event or (configuration?) recognizer for the syllable "pa" that occurs in that word. That ECS does not accept the sound of opening a door as matching its referent. And similarly the configuration ECS for the phonemic segment /p/.

Sensations and intensities are probably not specialized for language. But I understand that an ECS for a relationship among words or a word sequence does not ordinarily accept non-word perceptions as input, though the iconicity of reduplication may be an exception. Perhaps a word-class recognizer might also control for class or category membership of non-words, though I doubt it. The use of a program both for words and for non-word perceptions seems even more dubious. That is why I separated them. The question for me was, what is the mechanism for associating a word perception, or a word-dependency perception, with a nonverbal perception? How does the control hierarchy bridge the gap between the two sorts of perceptions.

You are telling me that my presupposition was wrong, that there is no gap to be bridged:

>that is not a
>physical division, only a division of convenience, a way of categorizing.
>This division is misleading, because it implies that there are
>perceptual, comparison, and output functions concerned EXCLUSIVELY with
>words, so that one could tease the two hierarchies apart into physically
>separate strings of nerves with some undefined sort of connection running
>between them. This is not my picture of the brain.

>What makes a perception into a "word" is not anything special about that
>perception, but the way it is used, the context within which it is used,
>the higher-level systems that use it (my version of Wittgenstein's
>principle). If word sequences are significant, they are significant to
>the same level of sequence recognizer-controller that deals with all
>other kinds of sequences -- without it, word-sequences would be
>unrecognized in the dimension of sequence, and they could not be
>produced, either.

Help me understand you. It seems to me that this second paragraph says nothing more than what I have said. A word-recognizer is just like any other event-recognizer in itself. However, its connections to other ECSs on higher and lower levels (and perceptions associated with it in memory?) make it unique.

To this I was adding that some ECSs have more in common in their connectivity and context of use. Looking at an imagined model, we from the outside can say that they cluster into different kinds, though the clustering may not be physical. Is this incorrect?

I am troubled about the role of language in hierarchical control and by the capacity of language to confuse us. I am troubled because we seem to be dependent upon shared language for representing our perceptions to one another. For these reasons, looking at the imagined model from the outside, the distinction between a kind of ECS associated with language and those that are not is important to me. I don't think that you are saying that all ECSs are associated with language.

I am not saying that there must be structures or "brain organs"

specialized for language. I think it is likely that ECSs involved in
language control are to a degree distributed among ECSs for nonverbal
perceptual control. I don't think the data on aphasia and brain damage
support the radical extreme of this view, that e.g. vocabulary for
locomotion is collocated with ECSs for locomotion. But I don't know
this. Is that what you are saying?

>long ago. For example, in your insistence that the structure of language
>is a social convention you were telling me that the structure of language
>that you were investigating is NOT a property of the human brain. To the
>extent that language is conventional, its structure is optional. Only
>those aspects of it that are exactly the same for every single human
>being can reflect true properties of and basic functions of the brain.

Some aspects of language structure appear to be universal and to reflect
true properties and basic functions of the brain. Others vary (are
optional) and appear to be conventional. The theory of language I have
been discussing is concerned with both, and makes the distinction on a
well defined basis and I think in a way that is useful to you.

The relation between the two is important for the vexed question of
social influence.

Bill Powers (920131.0800) --

>what is gained by using such awkward and ambiguous expressions?

Awkward, yes. Ambiguous, no.

The purpose is to have an explicit 1-1 correlation of form with meaning.

In most cases, the reductions (taken singly) are optional do not involve
any artificiality. Thus, I zeroed "the purpose of using such
expressions" to "the purpose" in the preceding sentence. The former is
more wordy, and if I wrote like that a lot it would be harder to
understand because of all the repetion. Probably a matter of short-term
memory, tuning out the familiar, and losing track of when to tune in
again for new information, or that's my guess.

For constructions that are relatively new developments in the language,
such as the progressive use of -ing, proposed sources and sometimes
intermediate steps can be quite awkward. This is the same thing as
saying that the reduction is obligatory (or nearly so) rather than
optional.

Bill Powers (920201.0900) --

>In your illustrations of the method of Harris, certain phrases come up
>that are considered by many people, perhaps even every person who speaks
>English at all well, as unacceptable. Can you think of any *a priori*
>reason why people would consider the following phrases unacceptable, even
>if they weren't being compared with other phrases?

>The slipper chewed the puppy
>The ball ate the ice cream
>The ice cream ate the ball

>I believe that Harris would say these phrases are unacceptable because in
>each case the subject of the phrase is an operator that does not take an

>argument, and so can't be connected by a transitive verb to an object.
>That isn't an a priori reason, but a generalization derived from
>considering many comparisons of phrases by individuals. The explanation
>should be stated more empirically: it has been found that ...
>
>The *a priori* reason I'm thinking of is simply that the implied events,
>which we can imagine, never actually happen. We have never seen anything

No, Harris would not predict the unacceptability of these sentences from the operator-grammar description of them. Just the reverse, in a sense. He would say that some of the words are operators, that certain operator-argument dependency relations are there, and certain reductions, because differences of acceptability (in the judgment of native speakers of English) are preserved when you compare one set of sentences (perhaps including some of these) with other slightly different sentences. The judgement of acceptability is a datum, a black box into which Harris does not peer. From the ranking of acceptabilities he can verify that certain sentence-differences involve only asserting additional operators or carrying out reductions, and others do not. He does not predict the unacceptability; he uses the preservation of acceptability-differences as a criterion to verify proposals about the structure of the language.

Of *course* the differences in acceptability are due to our ability to imagine circumstances of which we might say a given sentence. "The slipper chewed the puppy" might be said in context of describing a dream or a cartoon, "The Slipper's Revenge." Fine. If some other set of sentences is related by operators and reductions to the one containing this sentence, then the corresponding sentence ("it was the slipper that chewed the puppy" or "the slipper began to chew the puppy" maybe) has the same restriction as to the context in which it makes sense. The details of remembered and imagined perception, being an intractible morass, are tidily encapsulated in this simple criterion. This is what I mean by a constrained appeal to meaning.

>   By this method, Harris
>is not really exploring the world of language. He is exploring the world
>of perception.

I agree that the conclusions Harris reaches about language can tell us a great deal about the world of perception.

>If you were to make up another set of sentences relating to aspects of
>the world of which a person has had no experience, judgements of
>acceptability would mean nothing:

Yes, I agree. This becomes important for sublanguage grammars. It is why the changing structure of the sublanguage of immunology described in _The Form of Information in Science_ reflected the concurrently changing understanding of the field. Before a certain development, immunologists were "naive" about T-cells, after it they were not. A naive person would say "antibodies reject the host" or "the resistor heated the current". A person who knew something more about immunology would not say the first, and a person who knew something about electronics would not say the second. Here, "naive" is completely coincident with not knowing the sublanguage, and "knowing" the field is completely coincident with knowing its sublanguage. One cannot learn the sublanguage without controlling the perceptions of the field, and conversely one cannot learn to control the perceptions proper to the field without learning its sublanguage vocabulary and grammar. Resistance to change is perhaps often mostly resistance to sublanguage change.

>I repeat: this is just asking how the world of perception
>works, with the words being used merely to indicate the perceptions.

It is tricky when expectations (memory and imagination) arising from association of words swamp ongoing perceptual input. Happens all the time, right?

>These partial orderings Harris takes as properties of
>language itself.

He does not make that claim.

>If the same partial orderings are found in another
>language, Harris assumes that they indicate invariants of language that
>are characteristic of human Language.

He says nothing about the partial orderings beyond their availability for verifying details of the grammar, one at a time.

>This being an empirical study, the
>question as to WHY these orderings appear is not asked. That is, Harris
>does not look to any other level of explanation of the orderings as a way
>of predicting whether one ordering will be observed while another will
>not.

You are inditing Harris for not knowing about control theory. Sua culpa. Let's make up for the omission, shall we?

(Martin Taylor 920203 11:45) --

The two senses of "asbestos" (distinguished as chrysotile and crocodilite only in a mineralogy sublanguage) echo some of Whorf's experience with insurance claims. One example, people read "inflammable" on the sides of "empty" barrels as "not flammable," with explosive and expensive consequences. I'll look for the March 2, 1990 of Science.

There's an article "Inside the mind of a monkey" in _New Scientist_ for 4 January 1992. There have been several lately on whether or not simians have "self-awareness" required to attribute mental states to another. The claim is that they are good behaviorists but not good psychologists!

Also, an article on schooling in invertebrates, in NS for 23 March 1991. And an article in yesterday's Boston Globe described use of infra-sound (very low frequencies) by large land animals, mostly large mammals, but also crocodiles. These signals are quite complex, and carry for many miles over land, though of course not the 10K mile range of whale signals in the ocean (before the clutter of ships' engine noises).

Greg--I am most impressed with the span of time from which you drew material for the "social control" issue. I can imagine that in a way making it topical makes it more manageable. That's still a heck of a lot of ground to cover. Kudos!

Bruce Nevin
bn@bbn.com

Date: Thu Feb 06, 1992 11:07 am PST
Subject: Beer's Bug; PCT and PCT

From Pat & Greg Williams (920206)

>Bill Powers (920205.0800)

>To make a slippery surface, you would have to allow the feet to
>move backward without moving the body forward. And then, because the
>output is a force, you'd have to add the way leg position depends on
>neural output when resistance to leg movement is reduced.

Yes. Sorry we forgot to add that requirement.

>Maybe we could bypass all this guessing if you would just post the
>calculation that relates body velocity to leg force.

Current forward body velocity is simply proportional to current total backward force provided by all legs with feet currently down ("stuck" to the ground in Beer's model, or with that condition relaxed for slipperiness).

>Ooh, ow! Just remember that the neural circuits Beer omitted are those
>that didn't seem important from a non-CT point of view.

That doesn't follow. Beer's model bug controls in ways different from your own (upper alters lower reference signals) how-control-works model, but it still controls. Your problem seems to be with open-loop (see, we didn't call it "feedforward") at ANY level, but closed-loop at EVERY level just doesn't seem necessary for global control. An organism can "simply respond" (no resettable local reference signal) and then correct after-the-fact, sort of like a sampled-data control system, or perhaps not correct at all (if it "doesn't really matter" to survival, or, in a less ultimate way, to higher loops).

>If we can catch Beer's model in a wrong prediction, then it doesn't matter if
>the circuits he presented correspond to (selected) neural circuits: he's got
>it wrong. So far it looks as if his model predicts incorrectly when we
>introduce real physics into it. Unless cockroaches behave like his model
>would. It would be nice if Beer had based his model at least in part on
>behavioral data under appropriate external disturbances.

Do we detect a double standard here? Be careful to apply the same criteria to models made by whomever. We needn't point out the current infancy of PCT models. By "lesion" studies on his model, Beer has shown that under conditions of certain legs being disabled, the gait patterns correspond remarkably well to experimentally observed gait patterns. Do you, for example, know how well the Little Man's arm movements in tracking correspond to human tracking movements? I have some limited data indicating that there is not very good correspondence in at least one class of tracking movements. That doesn't mean I want to simply say "he's got it wrong" and start anew. All models can be improved -- the interesting question is when they are good enough for the purposes at hand. And both Beer's bug and your Little Man are good enough right now, we think, for some important purposes.

>The reason I'm not hot about neural modeling is that it's just analog
>computing done with elements that purport to represent neurons. Calling

>the elements neurons doesn't change the computations.

Certainly. But the idea is to incrementally improve the computations so they match more and more of the experimental data, rather than go off speculating wildly in realms where little or no data exist. AND the use of "neurons" as the basic building blocks, as opposed to building blocks at higher levels of integration, taking underlying processes as givens, requires filling in the details and not being able to postulate them -- sort of like, in programming, being given a high-level spec and having to write the low-level routines to accomplish it. No "miracles" allowed!

>The connectivity of the real nervous system is certainly suggestive of the
>system's gross organization, and the idea of neural oscillators is useful, but
>the actual representations of neurons and the functions they compute is very
>limited.

True, but future models should be less limited. This incremental approach to building more "accurate" models might be frustrating to some, but it is required unless we get the laws -- and the details -- of nervous systems handed to us by some Moses.

>I think we could have come up with the concept of neural oscillators
>without knowing the circuitry, and that we would have guessed that
>sensing positions of the legs would be important in creating finite back-
>and-forth movements. Even just guessing, we probably could come up with a
>walking bug that works quite like Beer's model.

The problem, as we see it, is in ONLY guesssing, and not paying sufficient attention to the data. Modeling alone is fine for heuristic purposes, particularly when "getting started" -- but to choose between alternatives (sometimes as different as day and night) requires listening to what the organisms themselves "say" in experiments.

>If Beer had an exhaustive model of all the neural connections, all the
>sensors, and all the functions computed by each neuron, he would have a model
>built from first principles. But I don't think that's what he has.

No one else has such a model, either, at this stage. The point is that his model is heading in a data-bolstered direction which WE don't think is contrary to the basic tenets of psychological control theory, although it might contradict some of the details in your proposed mechanisms for how control happens. Resolving that contradiction -- from both sides -- is going to take more than faith in principles somewhat removed from the data.

Pat & Greg

Date: Thu Feb 06, 1992 11:21 am PST
Subject: Misc comments

[Fron Bill Powers (920206.0900)]

Wayne Hershberger (920204.2131) --

> Configurations do appear to serve as sacadic targets. Try counting the
>number of letters in the rows below by saccading from one to the next;
>the first row is much less difficult.

Nifty example. When the configurations are all the same you lose track of where you are. I wonder -- is this because we saccade to configurations, or because we saccade with some position error? If you're looking at an "a" you can see that the next configuration is an "s" without moving the eyes, because both letters are within the 2-degree foveal region (1/10 inch separation at screen viewing distance of 20 inches = 1/200 radian or about 1/3 degree). So after the saccade, the fixated letter should be the one that was to the right. If there's a saccade error of 2/3 degree or more, you'll see the wrong letter, and know it, and do another saccade. But in the row lllll... you don't get any configuration error and don't know you've skipped one or more.

The actual target for the saccade could be at the sensation level, with the configuration check being done only after the saccade.

I was thinking in terms of separations of several degrees or more. In the following,

```
     v        r          s           u          n         t
```

... when I look at the v from 20 inches, the r is slightly recognizeable, and the s is problematic. Looking at the v, I can see all the letters, but not as letters. From u through t the target could be any letter. It's just a sensation of something.

When the letters are bigger

```
    V        R          S            U          N         T
```

I don't seem to do any better at recognizing the off-axis shapes. Maybe they just aren't big enough, so off-axis visual acuity is involved, but maybe configuration-recognition is confined to the central region of actual looking (not "attending" as I first proposed). This turns out to be a rather nice demonstration of the difference between sensation-level perception and configuration-level perception.

>There is another problem with representationalism; it is solipsistic.

This is not a problem with representationalism; it is a problem with language. You claim that "that to which the term representationalism refers" (for you) has the character of "that to which the term solipsism refers" (for you). Arguments at this level of verbal abstraction have little shareable meaning, as you and I have found out.

>>To the bug, the signals are the reality -- just as they are for us.
>Oh? What happened to the boss?

The boss reality is hypothetical. Hypotheses can be tested, but never wholly verified. They can, of course, be disconfirmed. The question of their validity can't be settled by verbal reasoning, but only by experiment.

>To me, the organism-environment interface is every bit as artificial
>(nasty) as the mechanism-environment interface.

I have no problem in distinguishing between the part of my model that refers to an organism and the part that refers to its environment, and between the part of the organism that's the CNS and the part that's topologically outside the CNS. The interface, in the latter distinction,

is easy to identify: sensory receptors and motor (or other) nerve-endings. If you don't make this distinction, you can't make a working model. Of course you can adopt a highly abstract point of view and say "it's all just matter and energy, so there's no difference." From that point of view, there is no difference. But then there's no model, either.

Bruce Nevin (920205.0656) --

The "I cut, you choose" scenario is very interesting. It allows the two sides to control for different variables that require the other side's cooperation -- but avoids the issue of who's controlling whom. When there is hatred and mistrust between the sides, each side not only wants to get what it wants, but it wants the other side NOT to get what it wants, the dog-in-the-manger syndrome. If it's agreed that the real wants can be concealed (the values which represent the real reasons for wanting the conditions being traded), neither side can maneuver to frustrated the other's higher-level objectives. I think that to reach this kind of agreement, there has to be an absence of hatred and mistrust. Maybe that's why it hasn't been tried yet.

I've heard it said that the best contract is the one that leaves both sides feeling they got the better bargain.

Joel Judd (920205.0938) --

RE reorganization:

>the less constrained the blind variation, the more variety in learning
>"outcome."

This isn't what I meant, but it's a good observation if we add a little more detail.

If there's only one outcome that will correct the intrinsic error (the Skinner box), then reorganization will converge to that outcome if it's found in time to prevent death. So the variety of outcomes is very small when the possible means of achieving an outcome are few. But if there are many alternative means that will create the same outcome, the same reduction of intrinsic error, then reorganization will cease when the first one is found. Reorganization being random, you can't predict WHICH means will end up being used, so there will be variety from one organism to another.

Reorganization itself is completely random (that is, it doesn't need any systematic component). But it acts within physical constraints. For example, in Martin Taylor's experiment with reversals, reorganization of imagined depth information at critical places on a visual object (my hypothesis, not his) can only work at those critical places, which are defined by the lower-level perceptual functions, the object being visualized, and the functions that derive 3-D information from 2-D displays. E. coli varies its direction of swimming randomly, but only in three dimensions of linear movement -- it can't vary the curvature of its path or its speed while swimming. The rat in the Skinner box can vary the way it sits on or the leg with which it pushes on the lever, but only a push of the lever will deliver food and reduce intrinsic error.

In general, the conformation of the environment, the present organization of the nervous system, and the site of reorganization determine the possibilities among which a random reorganizer can randomly choose. I

think this is important in development. The nervous system, when it begins working, has some initial organization handed down from previous generations. This organization provides all the possibilities that there are for the reorganizing system. I have some doubt that human beings would develop hierarchies of control organized around common classes of perception if evolution had not provided raw material making these classes the playing field within which reorganization is free to create specific examples. My hunch is that creation of so much similarity would be impossible if the new CNS were a total tabula rasa.

Date: Thu Feb 06, 1992 11:42 am PST
Subject: Beer Visit

[from Gary A. Cziko 920205]

Randy Beer (of Bug fame) gave a very interesting presentation yesterday evening on this campus. He presented his computer simulated cockroach, showed a video of the robot implementation of his model, discussed some work using neural nets and back propogation to develop circuitry for the flight response of the cockroach, and presented some results on the use of genetic algorithms (GAs) to artificially evolve neural ciruitry for leg controllers.

For me the most interesting part was the discussion of GAs to develop neural circuitry. Two questions came to mind: (1) was the environment in which he was evaluating the fitness function of his circuits "real" enough (with disturbances) so that control systems would develop; (2) could PCT/HCT modellers find this method useful for evolve hierarchies of control systems selected for their ability to control higher- and higher-level variables.

I also had a chat with him this afternoon, particularly on the role of what he calls "central pattern generators" in his model. His feeling is that real bugs have central pattern generators that provide command-driven behavior that are modified by sensory information. In his models, he tries to have the "best" of both worlds so that if there is no sensory information available (or things are happening to fast for it to be of use), the central pattern generator can nonetheless keep the behavior happening (he referred to the escape behavior of the cockroach as "essentially ballistic"), albeit in a rather crude fashion, but when sensory information is available it can also be used. He doesn't see why it has to be all one way or the other and that in fact both types of processes occur and interact.

I very much enjoyed his presentation and my chat with him. Randy is a very likeable, energetic young fellow with quite impressive knowledge of computing, biology, and neurology. I found it very interesting to contrast his "eclectic" approach with the "pure" HCT approach. Perhaps one day I'll understand them both better to make some better-informed judgements.--Gary

P.S. Bill, the video showed that the robot seems to act exactly the same way on its back as on its feet, except that it doesn't move forward when on its back (but then neither do real bugs).
==
Gary A. Cziko Telephone: (217) 333-4382

Date: Thu Feb 06, 1992 2:38 pm PST
Subject: Reorganization and Variation

[from Gary Cziko 920206.1000]

Bill Powers (920206.0900) said:

>If there's only one outcome that will correct the intrinsic error (the
>Skinner box), then reorganization will converge to that outcome if it's
>found in time to prevent death. So the variety of outcomes is very small
>when the possible means of achieving an outcome are few. But if there are
>many alternative means that will create the same outcome, the same
>reduction of intrinsic error, then reorganization will cease when the
>first one is found. Reorganization being random, you can't predict WHICH
>means will end up being used, so there will be variety from one organism
>to another.

I think this is a very useful way of looking at reorganization and can lead to useful predictions concerning the variability, especially cultural variation among humans. Humans have so many ways of satisfying their goals, and certain societies afford more variation than others. Other animals seem to have less opportunities and so we see more similarity in their behavior.

But it may lead to problems if applied in educational and other social settings. This is because it suggests that to get someone to reorganize the way *I* want this person to reorganize, I should control the environment so that only one way (the "right" way) can lead to satisfying goals. So, for example, I wouldn't let my daughter go out with anyone who is not white, educated and at least middle class.

But then, of course, I may be causing higher-level errors in my daughter, errors relating to her goals for independence in choosing her friends.

Oops, looks like we're slipping back into social control again. Sorry about that.--Gary

Date: Thu Feb 06, 1992 3:10 pm PST
Subject: A recent publication in the American Sociological Review 12/1991

Dear CSGers,   ++++FROM CHUCK TUCKER 920206++++

The ASR recently contained an article by Peter J. Burke "Identity Processes and Social Stress" where the abstract reads as follows: "Social stress can be understood by incorporating interruption theory as developed in research on stress into a model of identity processes drawn from identity theory. From this perspective, social stress results from interruption of the feedback loop that maintains identity process. I discuss four mechanisms of interruption of identity processes: broken identity loops, ... over-controlled identity systems ...." The article explicitly mentions Powers and has a circular diagram showing the closed loop but unfortunately the author understands little about our model and about feedback in the sense we use it but this might be a foot in the door. Regards, Chuck

Date: Thu Feb 06, 1992 3:38 pm PST
Subject: Use Your Foot, Tucker

[from Gary Cziko 920206.1122]

CHUCK TUCKER 920206 said:

>The ASR recently contained an article by Peter J. Burke "Identity Processes
>and Social Stress" where the abstract reads as follows: "Social stress can
>be understood by incorporating interruption theory as developed in research on
>stress into a model of identity processes drawn from identity theory. From
>this perspective, social stress results from interruption of the feedback
>loop that maintains identity process. I discuss four mechanisms of interrup-
>tion of identity processes: broken identity loops, ... over-controlled
>identity systems ...." The article explicitly mentions Powers and has a
>circular diagram showing the closed loop but unfortunately the author
>understands little about our model and about feedback in the sense we use it
>but this might be a foot in the door.

I assume that ASR is the BIG journal in sociology. So why don't you (and/or McPhail and/or McCLelland) use this opportunity to present the "right" way of using PCT in sociology by writing a commentary on Burke's article?

This reminds me to ask Bill Powers what happened to the response he was preparing for Bizzi's article which appeared in _Scienc_ I while back. I've been looking each week so I don't thing I missed it.--Gary

Date: Thu Feb 06, 1992 6:07 pm PST
Subject: Interfaces

[from Wayne Hershberger]

(Bill Powers 920206)
>maybe configuration-recognition is confined to the central region
of actual looking

Yes, that's where I'd put my money.

>I have no problem in distinguishing between the part of my model
>that refers to an organism and the part that refers to its
>environment, and between the part of the organism that's the CNS
>and the part that's topologically outside the CNS.

Once one has delineated all the various and sundry parts comprising the whole model, the environmental parts may be distinguished from the organismic parts with a "dotted line" drawn according to established conventions--no problem. The problem to which I think Joel Judd was alluding concerns the interface between the whole model and its metaphysical context (boss reality). Where does one draw THAT dotted line?

Warm regards, Wayne

Wayne A. Hershberger   Work: (815) 753-7097
Professor of Psychology

Department of Psychology   Home: (815) 758-3747
Northern Illinois University
DeKalb IL 60115   Bitnet: tj0wah1@niu

Date: Thu Feb 06, 1992 6:21 pm PST
Subject: Misc comments

[From Rick Marken (920206)]

Won't anyone tell me what "cognitive penetrability" is? Or do I have to go to the library?

Bill Silvert (920205) says:

>The amount of traffic on this mailing list is oppressive enough.

How does mail traffic oppress? You seem to feel that you are being pushed around by an external stimulus (the mail traffic). As far as I can tell, mail traffic can't make you do or feel anything unless you have goals with respect to it. Although the traffic itself cannot be controlled, the amount you deal with can; I find that pressing "Control C" works for me. Feeling "oppressed" by a perceptual variable is a sure sign of internal conflict. It seems to me that you might have two conflicting goals with respect to mail traffic 1) you want to read it and 2) you want to spent a short time doing it. I suggest that "going up a level" might help suggest ways to set these goals so that they are not in conflict. Then you can feel "in control" of the mail traffic, rather than vice versa. I think this is a better approach than trying to control the control systems that are the source of the mail traffic.

Gary A. Cziko (920205) on Beer:

>   In his models, he tries
>to have the "best" of both worlds so that if there is no sensory
>information available (or things are happening to fast for it to be of
>use), the central pattern generator can nonetheless keep the behavior
>happening (he referred to the escape behavior of the cockroach as
>"essentially ballistic"), albeit in a rather crude fashion, but when
>sensory information is available it can also be used. He doesn't see why
>it has to be all one way or the other and that in fact both types of
>processes occur and interact.

How does the nervous system know when to be one way and when to be another? Is that Beer's contribution -- has he discovered how the nervous system knows when to switch from being open loop to being closed loop? What is "command driven behavior modified by sensory info"? Aren't the commands then reference signals -- specifying intended states of perceptions?

It has to be one way or the other because it can only work one way -- closed loop. If the organism controls in an unpredictably variable environment then it MUST be organized as a perceptual control system (this was Bill's point in the "Behaviorism" post included with the last newsletter). There can be "open loop" components of a control system (as Greg and Pat note in their recent reply to Bill) , but the variable that is the "output" of such a open loop component is NOT controlled. For example, here is "open loop control" of muscle tension by error signal -- but that's not "control", it's cause. There could be several open-loop segments in the control loop, but none of them control.

This whole thing with Beer (and Brooks and the others) is driving me relatively nuts. What is their big contribution to our understanding of living systems????? What are they doing that hasn't been done in a variety of incarnations since, at least, the seventeenth century. I have no doubt that Brooks, Beer, etc can make toys that do very impressive things-- heck, German clockmakers made some pretty impressive automata over 300 yrs ago. What new priciples have Brooks and Beer and whomever discovered? What basic insights about the nature of living systems have they discovered? Central pattern generators? -- give me a break! Using feedback until it starts to happen too fast? This is kindergarten stuff.

I think that the discussions of the various clever mechanisms that might be used to produce interesting "behavior" are obscuring a basic question -- is behavior a process of controlling variables?? That is, is behavior an example of the Phenomenon of Control. Beer, I believe, would say (to the extent that he understood the question) "only partly. Much behavior is just generated output." That is why his models can be partly s-r and partly "modified by sensory feedback". PCT is only interested in behavior that involves control -- because that is what it is designed to handle.

It is true that Beer's bugs control with respect to a fixed reference at low levels -- because of the fact that SR systems with appropriate dynamics are equivalent to control systems with fixed references. But Beer misses the point that behavior is the control of perceptual variables; the bug is not controlling it's leg; it is controlling the sensed angle of the leg with respect to the body, a variable that is influenced by the position of the leg and possibly many other factors -- many of which are not influenced by the organism. This is the important organizational principle that comes out of PCT; Behavior (that is, intentional, purposeful results of action) is Control -- and what is controlled is PERCEPTION. Obviously, this principle has enormous implications for our understanding of the nature of living systems -- but you would never know about it from looking at Beer's models; even if this principle is embodied in some aspects of the model.

As far as "disproof" of Beer's model. How about using his model to produce the behavior of a person making a rhythmic circle on a computer screen -- with the cursor under disturbance? Could someone show me a diagram of Beer's model of that? I bet it would have to be a control model -- the "central pattern generator" for circular motion of the cursor would have to be a reference for the perception of that movement -- wouldn't it?

Best regards

Rick

Date: Fri Feb 07, 1992 1:42 am PST
Subject: Re: Beer's Bug; PCT and PCT

[Oded Maler 92.02.06]

Beer's thesis/book was written in the period of the new revival of neural network, and this is reflected in the aspects that are over-emphasized in his work, namely the fact that some functions from input to output can indeed be computed by neural-like circuitry. Too much, to my taste, is devoted to the simulation of the neural

nets, and too little is devoted to the dynamics of the interaction
with the physical environment, which is done mostly on
qualitative-level. This tacitly assumes, either that the world
is disturbance-free, or that lower-level control systems already take
care of this disturbances, and that the higher-level coordination
system can thus assume a disturbance-free environment.

Thus, what is demonstrated is essentially that 1) certain neural-like
circuits can evoke certain kinds of 6-dimensional sequences of {0,1}
that correspond to walking, and that 2) connecting these circuits to
abstract touch sensors will produce a turning and 'retract' behavior.
(The correspondence to neurbiology is a bit speculative - like in many
other works at the current state of knowledge)
These are nice results but not as revolutinary as one might imagine
from the hype - but this is how science works these days.

For those interested, a good survey of biological models and data
concerning these topics is: H. Cruse, Coordination of leg Movements in
Walking Animals, in Meyer & Wilson (Eds.), Simulation of Adaptive
Behavior, MIT Press, 1991. It contains a lot of references.
Cruse's e-mail is UBIOF140@DBIUNI11.bitnet

--Oded

Date: Fri Feb 07, 1992 7:21 am PST
Subject: Beer's Bug

From Pat & Greg Williams (920206)

>Gary A. Cziko 920205

>P.S. Bill, the video showed that the robot seems to act exactly the same
>way on its back as on its feet, except that it doesn't move forward when on
>its back (but then neither do real bugs).

We're surprised. Probably the influence of the forward and backward angle
sensors on the gait oscillators is not as great as we thought.

>Rick Marken (920206)

>Is that Beer's contribution -- has he discovered how the nervous system
>knows when to switch from being open loop to being closed loop?

It doesn't necessarily have to switch. One of the "modes" of behavior of
Beer's bug is "edge-following." What happens is when an antenna hits a wall,
it CAUSES the bug to turn away from the wall and also activates a slower-
acting circuit which turns the bug back into the wall (meanwhile having moved
forward), and the process repeats. Thus, the bug "bounces" its way along a
straight edge, via a process incorporating both locally open-loop (no
adjustment with respect to perceptual changes) AND closed-loop (the "kick"
away from the wall depends on the antenna signal) processes which are
happening in part simultaneously.

>What is "command driven behavior modified by sensory info"? Aren't the
>commands then reference signals -- specifying intended states of perceptions?

Better ask Beer himself to make sure, but we think he means that "command driven behavior" is locally open-loop and proceeds without perceptual signal/reference signal comparison, but every once in a while there could be corrections made to that behavior which DO involve p.s./r.s. comparisons. In other words, there would be no error-correction for a while, then correction. Or there might be no correction under some environmental conditions, correction under others. The "commands" would be specifications for output patterns, not input patterns, and those output patterns would be subject to distortions due to disturbances in between periodic (if ever) corrections.

>There could be several open-loop segments in the control loop, but none of
>them control.

We agree. But it appears that the Powers control theory mechanism (not the underlying notion that organisms control) denies that there can be local open loops which don't get corrected CONTINUOUSLY. Yet we don't see much evidence to choose the Powers-type mechanism over the Beer-type mechanism. (No doubt there are other candidates -- it isn't either-or, necessarily.) Very little is known experimentally about the human control structure, especially at the higher levels. How those higher levels work is guesswork at present. Our plea is for an open mind among psychological control theorists, not for S-R models but for non-Powersian control mechanisms such as those in Beer's bug.

>This whole thing with Beer (and Brooks and the others) is driving me
>relatively nuts. What is their big contribution to our understanding of
>living systems????? What are they doing that hasn't been done in a
>variety of incarnations since, at least, the seventeenth century.

We think Beer is showing other-than-Powersian mechanisms for psychological control which are at least possible, as noted above.

>I think that the discussions of the various clever mechanisms that might be
>used to produce interesting "behavior" are obscuring a basic question --
>is behavior a process of controlling variables?? That is, is behavior an
>example of the Phenomenon of Control.

We have no quibbles at this level of discourse. Bill's insights about the need for control in environments with disturbances aren't going to be displaced. Beer's bug demonstrates behavior as control, in our opinions, and Beer's own interpretations are a side issue -- or should be -- for control theorists.

>It is true that Beer's bugs control with respect to a fixed reference
>at low levels -- because of the fact that SR systems with appropriate
>dynamics are equivalent to control systems with fixed references.

Are you now claiming that Beer's bug is an S-R machine, and/or only that he thinks it is? Not long ago, you posted a claim to the effect that Beer's bug was a psychological control system, but it LOOKED SUPERFICIALLY like an S-R machine.

Warm (not hot) regards,

Pat & Greg

Date: Fri Feb 07, 1992 8:25 am PST

Subject: Marken on Beer's Bug

[from Gary Cziko 920206.2115]

(Pat and Greg: Please fill in the details for me)

An obviously agitated Rick Marken (920206) writes:

>This whole thing with Beer (and Brooks and the others) is driving me
>relatively [sic] nuts.

Let me offer some advice by paraphrasing a leading PCT theorist and modeller:

As far as I can tell, Beer and his bug can't make you do or feel anything unless you have goals with respect to them. Although you cannot control Beer and his bug, you can control how you deal with them. Pressing "Control C" might help. Feeling that you are being driven "relatively nuts" by a perceptual variable is a sure sign of internal conflict. It seems to me that you might have two conflicting goals with respect to Beer and his bug: 1) you want to read about them, and 2) you want to see your current model of behavior explain everything. I suggest that "going up a level" might help suggest ways to set these goals so that they are not in conflict. Then you can feel "in control" of Beer and his bug, rather than vice versa. I think this is better approach than trying to control the control systems that are the source of Beer and his bug and people who discuss Beer and his bug.

>How does the nervous system know when to be one way and when to be >another?

Beer found that the slower gaits of his simulated bug depended more on sensory feedback than the faster gaits. Once the bug got moving fast enough he could remove the sensory feedback connections and it would still keep moving. He did not set up his circuits to act this way, they just did.

>It has to be one way or the other because it can only work one way -- closed
>loop.

I find the logic of this sentence less than totally convincing.

>If the organism controls in an unpredictably variable environment
>then it MUST be organized as a perceptual control system (this was Bill's
>point in the "Behaviorism" post included with the last newsletter). There
>can be "open loop" components of a control system (as Greg and Pat note
>in their recent reply to Bill) , but the variable that is the "output" of
>such a open loop component is NOT controlled. For example, here is "open loop
>control" of muscle tension by error signal -- but that's not "control", it's
>cause. There could be several open-loop segments in the control loop, but none
>of them control.

Beer and the Williamses are arguing that you don't need continous control at the lowest level. For another example, when I ask the waitress for a cup of tea I am not in continuous control of my tea-getting. If she comes back with coffee I just send her back for the tea.

>What basic insights about the nature of living systems have they >discovered?
>Central pattern generators? -- give me a break! Using feedback until it
>starts to happen too fast? This is kindergarten stuff.

I suppose I went to the wrong kindergarten.

>I think that the discussions of the various clever mechanisms that might be
>used to produce interesting "behavior" are obscuring a basic question --
>is behavior a process of controlling variables?? That is, is behavior an
>example of the Phenomenon of Control. Beer, I believe, would say (to the
>extent that he understood the question) "only partly. Much behavior is just
>generated output." That is why his models can be partly s-r and partly
>"modified by sensory feedback". PCT is only interested in behavior that
>involves control -- because that is what it is designed to handle.

So if real organisms use uncontrolled (at the lower levels) behavior you are not interested. That's your right. But why rule out open loops at these lower levels? In fact, it would seem that we should start with the simplest models (S-R) and then add to them as necessary to make our models match behavior. I believe even Bill Powers has said something similar. Shouldn't we add only as much control as is actually needed than to put it everywhere by fiat?

>Behavior (that is, intentional, purposeful results of
>action) is Control -- and what is controlled is PERCEPTION. Obviously,
>this principle has enormous implications for our understanding of the nature
>of living systems -- but you would never know about it from looking at Beer's
>models; even if this principle is embodied in some aspects of the model.

Maybe that's what cognitive inpenetrability is all about. Perhaps real organisms are also cognitively inpenetrable. I think the Williamses have already made this point.

>As far as "disproof" of Beer's model. How about using his model to produce
>the behavior of a person making a rhythmic circle on a computer screen -- with
>the cursor under disturbance? Could someone show me a diagram of Beer's model
>of that? I bet it would have to be a control model -- the "central pattern
>generator" for circular motion of the cursor would have to be a reference for
>the perception of that movement -- wouldn't it?

Cockroaches don't do such things, especially not when walking fast. Instead, why not create a model that walks like a cockroach but somewhat clumsily and is not able to compensate for relatively unimportant variables. You might wind up with something like Beer's.

I really don't think I'm the one to really add much to this discussion, but I coudn't resist giving some good counselling advice to Rick and I would like to keep the discussion going, particularly between Marken, the Williamses and Powers. I think I have much to learn here. And thanks to Beer, his bug and the Williamses for making us grapple with this.--Gary

Date: Fri Feb 07, 1992 2:53 pm PST
Subject: Beer's Bug

[From a somewhat calmer Rick Marken (920207)]

Pat & Greg Williams (920206) ask;

>Are you now claiming that Beer's bug is an S-R machine, and/or only that he
>thinks it is?

The latter -- though it is partly an SR machine, as I understand it. I consider the central pattern generator possibly SR unless its output depends on an input variable that is influence by its output. To the extent that there are unsensed consequences of the bugs output that are considered aspects of the bug's behavior, it is SR. This is actually one of the problems with not using control theory -- it's rather hard to say what the bug is DOING without the notion of a controlled variable. Its the same problem that behavioral psychology has in general. Without the concept of control, the behavior of the bug can be defined as an variable consequence of its neural activity. The "best" definition of behavior, from this point of view, is the variable that correlates best with some "causal" variable (whether inside or outside the bug). In control theory, the only variable consequences of neural activity that are considered behavior are controlled variables. Beer has simply defined the behaviors that he wants to see the bug perform; moving forward, near an edge, whatever. Some of these variables may be controlled; some not. But for Beer they are just "outputs".

A control theory approach to modelling bug behavior would start by either 1) defining the variables that the bug is to control (this was Bill P.s approach in the arm demo -- he just guessed at the variables controlled when pointing at an object in space) or 2) finding out (using the test) what variables a bug actually controls when doing a particular activity -- then modelling it.

So now I've clarified my problem with Beer a bit for myself (if not for others). Beer's modelling is aimed at producing a bug that will produce particular behavioral outputs, whether these output are controlled or not. While his model does, indeed, control certain variables (I bet you could apply the test to the computer bug to determine what these are) it is not built in order to control these variables. Beer's bugs are based on a behavioristic view of behavior itself (the alliteration in this sentence, for example, was a behavior of mine, but it was not a controlled result; it happened by accident).

Gary Cziko (920206) says:

>seems to me that you might have two conflicting goals with respect to Beer
>and his bug: 1) you want to read about them, and 2) you want to see your
>current model of behavior explain everything.

Yeah -- caught me. Hopefully what I said above suggests that horn # 2 of my dilemma is a little deeper than hoping that control theory will explain everything. My "going up a level" suggests that my conflict results from principles I have about dealing with first principles. The first principle that I think is being ignored in Beer's work (and other robotics work) is the question of "what are we trying to model?" As a psychologist, I think of the problem as "what is behavior?" I refer you to chapter 1 in my soon to be published Opus, for my attempts at an answer. If Beer, Brooks, et al thought more about these first principles they might devote their considerable skills to building models of systems that control perceptual variables --rather than models of systems that produce particular output variables.

>   But why rule out open loops at
>these lower levels? In fact, it would seem that we should start with the
>simplest models (S-R) and then add to them as necessary to make our models
>match behavior.

Ah Ha -- Match behavior!!! But WHAT IS BEHAVIOR. What are we trying to match. I agree with starting simple (god knows) before going to complex. But before we even start, we should decide what we are trying to explain. If "behavior" is "output" then I can try to mimic any variable consequence of action that I want. So "angle of arm with respect to ground" is a behavior. We can model this very easily with an open loop system that changes arm angle. But if behavior is "controlled perceptual variables" then we don't even need to consider open loop models -- we already know thay don't control.

I get a bit excited about this stuff, I think, partly because I think that PCT is the most incredible insight about the nature of living systems that has occurred in the last 200 years. It is also because I hate to see very skillful, clever, intelligent people, who are trying to understand living systems, miss the boat. I just wish someone with Beer's skills would start doing his work from a PCT perspective. As it is, I see great abilities being squandered on old misconceptions. With people like Beer or Brooks et al working from a PCT perspective, we could move much more quickly toward a more detailed understanding of the behavior of living systems. I'm sorry that it's not happening. C'est la vie.

Best Regards

Rick M.

Date: Fri Feb 07, 1992 6:03 pm PST
Subject: Language; perception; BEERBUG

[From Bill Powers (920206.1300)]

Bruce Nevin (920205.1303) --

>Are you saying that there are no program perceptions for syntax, for the
>conventions by which words are arranged in linear order? Or are you
>saying that those same programs apply equally well to non-word
>perceptions?

I don't know. I proposed a way in which words and meanings can be substituted at the input of a perceptual function, and I probably overgeneralized. Maybe what I was thinking of would apply best during communication, where it's the relation between words and meanings, not syntax that counts most (who really cares if I say whom?).

>Briefly, it appears to me that some higher-level control applies to
>words but not to their "referents" (nonverbal perceptions). Can you see
>a way past this?

Yes, I mean, no I can't. This weakens my "all one hierarchy" claim. If a perceptual function deals with words without meanings, that is, purely linguistically, it treats a word itself as a meaning to which other words (like "noun" or "operator") could point. So input functions designed to treat words themselves as input objects would deal only with words, relationships among words, categories of relationships, sequences of categories, and programs of sequences -- syntax. As you imply. So when we're communicating, we may treat words and meaning as equivalent in the process of communicating meanings, but when we're controlling for the structure of language we treat the basic words themselves as meanings and speak in a metalanguage about them as if they were just perceptions called "nouns," "verbs," "operators," and so on.

Is this getting us somewhere?

>"Slipper" is not an operator. The difference between the two sentences
>is that the arguments of the operator "chew" are in the reverse order:
>
```
>             N      Onn         N
>       The puppy   chewed the slipper.
>       The slipper chewed the puppy.
```

How do you know that "slipper" isn't an operator? Is there some characteristic of a word or its meaning that distinguishes operators? Is there some way you could distinguish an operator-term by the level of perception to which it refers? For example, can ANY sensation, intensity, or configuration be the meaning of an operator-word? Do all operators name classes of relationships? Transitions and events? Any ideas?

>... objection to the artificiality of some of his [Harris'
>sources and intermediate stages seems kind of silly. There is no
>objection to the metalanguage of semantic features being different from
>ordinary language.

I had understood the expanded forms to be those that even non-linguists would naturally produce when they want to be as explicit as possible. If you're saying that one can learn to construct such forms according to some formal algorithm, then I have to agree -- I have at least one demonstration before me. But I don't see any evidence that nonlinguists construct them. I don't knowingly use them myself. Are you saying that I use them without knowing that I do?

My present position (see date/time stamp) is that any "expansion" that goes on is done through imagining nonverbal perceptions to flesh out the scene that is suggested by the terse forms we do use. When we do verbally expand to be more explicit, I claim that we do so by describing the more complete scene or meaning that we have build nonverbally on the initial picture. We don't need this description to construct a satisfactory meaning from the original sentence. To elaborate on an earlier post:

Harris' proposal:

  Short sentence -> long sentence -> explicit meaning

My proposal:

  Short sentence -> sketchy meaning -> explicit meaning [-> long sentence, optional]

I don't have much confidence that I really understand Harris' propositions, even after having read a little of his work. Is the above contrast valid?

Martin Taylor (920205.1345) --

Re: perceiving both orientations of Necker cube simultaneously

>Here's the problem with introspection. I disagree, based purely on my
>subjective experience [that both can't be perceived at once].

I have known only one other person (Sam Randlett) who claims to be able to perceive both orientations of the Necker cube at the same time. I spent about 20 minutes trying to do it, yesterday, and while I could get the orientations to flip back and forth fairly rapidly (once per second or maybe a little better over a short period) I failed utterly to see them both at once. Either one face of the cube was in front of the opposite one, or it was behind it -- I couldn't get the simultaneous sense of "in front AND behind." In fact, I couldn't even IMAGINE doing so, which of course only illustrates the fact that we perceive imagined information using the same input functions we use for real-time information. If I can't perceive it I can't imagine it either.

If you can do this and I can't, then there is a considerable difference between our ways of constructing 3-D perceptions out of 2-D perceptions. In fact your brain may do it according to the model you're proposing (two complete sets of perceptions) and mine may do it according to the model I propose (adding missing depth information). That wouldn't be too surprising, except to those who think the brain is hard-wired at birth. It would be really lovely to think of a control task in which the two modes of perception could be tested. What could one control that would depend on being able to see both orientations of a binary reversal at the same time? If we could devise such an experiment, you would be able to control the dual variable and I wouldn't be able to. That would be a truly smashing experiment! What about a figure in which the reversed form is not symmetrical in depth with the other form? Is that possible?

Roger on the spellings.

>... another case in which perceptual conflict can occur. It is called
>the Stroop effect.

I'm familiar with it as the "Stroop Test." You write a long list of color-names, writing each name in a color other than the named color (so "RED," for example, is written in blue). Then you're supposed to go through the list naming the color of each word rather than reading the word. So much conflict is produced that this test is used in stress tests. Here the conflict isn't between perceptions directly, but between actions: to say the name of the color, or to say the printed word. Evidently the sight of the color (sensation) leads to imagined perception of the color-name, while at the same time perception of the printed name is present. The conflict concerns which word to say: you can't say both at once. Naming a color conflicts with reading aloud. But there's no problem in seeing that there is a color with one name while the letters spell a different name. If you don't try to say the color names, or if you spell out each word aloud, there's no difficulty.

I'm not trying to be picky. It just seems to me that the term "perceptual conflict" is ambiguous. It could mean that somehow one perceptual signal is opposing another perceptual signal directly, or it could mean that there is a conflict attributed to presence of the distinct -- and non-opposing -- signals. If the perceptions coexist, they can't affect each other any more: they have already been generated. If they're merged in any way, then only one signal results and there's only one perceptual signal. Conflict, it seems to me, arises only when the perceptions imply mutually-exclusive actions of some sort. Then it is produced by the systems that are trying to take those actions, not by the perceptions. This is "perceptual conflict" only in the sense that it is somehow associated with perception; the perceptions themselves simply exist. Whether they result in conflict doesn't depend on the perceptions, but on what higher systems try to do with them.

The example you give actually supports my proposition:

>Imagine having to press a button marked "animal" or one marked "tool"
>and being presented with a cartoon horse with the word "hammer" written
>across it, and being asked to press the button according to which
>picture was being shown.

If you're allowed to use both hands, there's no conflict: you press both buttons. So clearly there's no interaction between the perceptions. The conflict arises from the (implicit) condition that you can press only one button at a time. So one comparison says "press left" and the other says "press right;" in the first half ss down on the
table between the buttons, refuses to move, or goes into lateral oscillation. The perceptions themselves aren't in conflict: the goals for action are.

>There must be a place where the two sources of information converg
≡ì~e to
>produce the same result, if you have to make an action choice based on
>one or the other, but they are equally clearly separate in their effects
>in many ways.

Sources of information don't have effects on behaviors except through comparison with reference signals -- not in the CT model. And, considering the hierarchy of perception, the perceptual effect of a source depends only partly on the source; the rest of the effect depends on the level and on the particular interpretation at that level. Isn't your interpretation rather subtly invoking an S-R model?

Gary Cziko (920205.1557) --

Thanks for the report on the Beer talk. With respect to open-loop pattern generators versus "use" of sensory information,

>He doesn't see why it has to be all one way or the other and that in
>fact both types of processes occur and interact.

There's nothing that says open-loop systems can't exist. They can. The question is what they can accomplish. Maybe if a cockroach responds to a puff of air by running like hell in whatever direction it's headed, it will escape more often than run into the jaws of doom. So you get a very primitive feedback effect that's mediated by evolution rather than by a neural control system. If that's the best the cockroach can do, then so be it. But it's not a very interesting system, in that case. It doesn't tell us much about the behavior of higher organisms, which is almost entirely control behavior. It's possible to build an open-loop cockroach, and if that's sufficient to reproduce its behavior under all circumstances, then it's sufficient and we shouldn't apply control theory. Is it sufficient?

What bothers me about Beer's statements is that he's treating the problem qualitatively, whereas you can't understand the role of sensory feedback without considering quantitative relationships.

>... the video showed that the robot seems to act exactly the same way on
>its back as on its feet, except that it doesn't move forward when on its

>back (but then neither do real bugs).

If real bugs move their legs the same way when they're on their backs, we can rule out control of forward velocity. If they speed up their leg movements or push harder when their bodies are prevented from moving (or when their feet slip) we can guess that there's control, and try to identify the controlled variable and measure the loop gain. Since none of this has been done (formally), we should just hold in abeyance any claim that bugs are control systems. You have to show that there's control before there's any sense in applying control theory. Once we know whether there's control, we can decide on how to interpret neural hookups.

If real cockroaches push harder when a force is applied to retard their motion, there are control systems. In that case, if Beer's bug doesn't push harder when retarded, his model is wrong in some detail. Perhaps it would just need some minor modifications, or perhaps some assumptions about what is sensed need to be changed. Some behavioral data are needed. On the other hand, the model might match the behavior perfectly. Who knows?

Pat & Greg Williams (920206.0810)

>>... because the output is a force, you'd have to add the way leg
>>position depends on neural output when resistance to leg movement is
>>reduced.

>Yes. Sorry we forgot to add that requirement.

How would you add it to Beer's model?

>Beer's model bug controls in ways different from your own (upper alters
>lower reference signals) how-control-works model, but it still controls.

I'll check it out. Turns out that I won't be buying Beer's book -- it's $30 -- but I'll get it through interlibrary loan and copy the parts I will need for reference.

>Do we detect a double standard here? [About testing models]

Do you know any model that's been tested more extensively and at a more basic level against real behavioral data than the CT model?

>Do you, for example, know how well the Little Man's arm movements in
>tracking correspond to human tracking movements? I have some limited
>data indicating that there is not very good correspondence in at least
>one class of tracking movements.

Not very good compared with what?

>All models can be improved -- the interesting question is when they are
>good enough for the purposes at hand. And both Beer's bug and your
>Little Man are good enough right now, we think, for some important
>purposes.

Point for your side. Maybe three.

All your points are well taken. I'm only resisting the idea that neural circuit-tracing ALONE is sufficient to lead to an adequate model. It can certainly help to get us on the right track and distinguish between

otherwise equally-convincing models. But there's an enormous amount of room for fudging when the circuitry is known only in its gross features. By selective circuit-tracing, you can show that the human nervous system connects sensory inputs to motor outputs, and justify S-R theory. It's been done. Behavioral experiments have a way of telling you that no matter how correctly you've traced every connection, you still have something wrong because the real system doesn't behave the way the model does.

Best to all,

Bill P.

9202B CSGnet

Date: Sat Feb 08, 1992 5:04 am PST
Subject: Beer's Bug

From Pat & Greg Williams (920208)

TO BRUCE NEVIN: We need your USPS mailing address, not e-mail address!

>Oded Maler 92.02.06

>Thus, what is demonstrated is essentially that 1) certain neural-like
>circuits can evoke certain kinds of 6-dimensional sequences of {0,1}
>that correspond to walking, and that 2) connecting these circuits to
>abstract touch sensors will produce a turning and 'retract' behavior.

We would like to emphasize that ALL models of sensory receptors are filters which "throw away" information about some types of disturbances. The question is whether a particular model handles disturbances of interest in the situation being modeled, and whether the real sensors which are being modeled also do so. Our retinal cells don't handle the disturbance of being exposed to a nuclear flash very well, and so one's model for them shouldn't handle it very well. It is true that Beer's bug's sensory receptors are very crude. But his model is a beginning, only. Our expanded NSCK program will allow the user to specify more complex sensory transducer functions.

>a somewhat calmer Rick Marken (920207)

>... it's rather hard to say what the bug is DOING without the notion of a
>controlled variable. Its the same problem that behavioral psychology has in
>general. Without the concept of control, the behavior of the bug can be
>defined as an variable consequence of its neural activity.

We think it would make a highly publishable paper, given all of the interest in Beer's bug throughout the behavioral simulation community, for some talented psychological control theorist out there (hint, hint) to show how CT "makes sense" of Beer's bug's workings. Beer himself just might be open to a collaboration on same, after he has been convinced by the talented (and presumably patiently understanding) control theorist that there is a better way to understanding "adaptive" behavior. We're NOT being facetious! CT HAS to make a dent SOMEWHERE, and this could be it. Beer is, at the very least, close to CT thinking... a slight shove, and...

>I get a bit excited about this stuff, I think, partly because I think
>that PCT is the most incredible insight about the nature of living systems

>that has occurred in the last 200 years. It is also because I hate to see
>very skillful, clever, intelligent people, who are trying to understand
>living systems, miss the boat. I just wish someone with Beer's skills
>would start doing his work from a PCT perspective. As it is, I see great
>abilities being squandered on old misconceptions. With people like Beer
>or Brooks et al working from a PCT perspective, we could move much more
>quickly toward a more detailed understanding of the behavior of
>living systems. I'm sorry that it's not happening. C'est la vie.

We think YOU could help make it happen. Beer's bug is getting raves from at least some AIers because it controls. Of course, they don't understand that. They talk about "robustness." Well, PCT is THE theory of robustness! Somebody needs to show them why it "works," when so many other approaches haven't. You have the theory behind the up-and-coming state-of-the-art, and only you (and other PCTers) know it. How about letting the rest of the world in on it? (I know, you've tried and partially failed with the psychological community -- but this is a whole new ball park, and the PCT team (under different management) just hit a home run.)

>Bill Powers (920206.1300)

>What bothers me about Beer's statements is that he's treating the problem
>qualitatively, whereas you can't understand the role of sensory feedback
>without considering quantitative relationships.

But behind his qualitative statements lies a quantitative model which can be interpreted (and improved) by PCTers, as noted above.

BP>>... because the output is a force, you'd have to add the way leg
BP>>position depends on neural output when resistance to leg movement is
BP>>reduced.

P&GW>Yes. Sorry we forgot to add that requirement.

>How would you add it to Beer's model?

As the model stands, the feet are either up or down (fixed). To add a "slippery" condition, when a foot was down, its leg force would need to divide between contributing to forward movement of the bug's body and slipping that foot backward. Using Beer's primitive force = k * velocity "physics" one could model the contribution to forward body movement as before and make backward slipping velocity proportional to a "slipperiness coefficient" * force, with the constraint that the total force balances.

G&PW>>Do we detect a double standard here? [About testing models]

>Do you know any model that's been tested more extensively and at a more
>basic level against real behavioral data than the CT model?

Which CT model? Have you compared detailed trajectories of the Little Man's arm with measured human arm trajectories? In your tracking models with high correlations, I (Greg) suspect that the high correlations are quite insensitive to actual model details (though I don't think you or Tom have run sensitivity studies with varying models as well as parameters) -- the correlations are high because the tracking is close to perfect. If you made tracking more difficult, the simple model wouldn't work as well (which is why the human factors engineers have much more complicated models, which, yes, I would say have been more extensively tested than most CT models). Of course, I realize that the HF models aren't truly generative, because they include empirically derived "curve-fit" terms. That is a real challenge for PCTers: develop a GENERATIVE model of DIFFICULT human tracking.

Pat & Greg

Date: Sat Feb 08, 1992 11:46 am PST
Subject: Social control; BEERBUG

[From Bill Powers (920208.1000)]

Gary Cziko (920206.1000) --

>I think this is a very useful way of looking at reorganization ... But
>it may lead to problems if applied in educational and other social
>settings.
>... it suggests that to get someone to reorganize the way *I* want this
>person to reorganize, I should control the environment so that only one
>way (the "right" way) can lead to satisfying goals.

This is how it's done already. The basic method of "teaching" is to define the result that's demanded and to make getting what you want contingent on producing it, with essentially no advice on how to do that. If you don't produce, you don't get the grade or the pay or whatever it is you need. How you manage to satisfy the requirement is up to you and your reorganizing system. Skinner merely formalized operant conditioning; it was the main technique of education, religion, and industry long before Skinner arrived on the scene. Behind this technique there must be raw physical power, so the contingency can be established without the consent of the controlled, and so the reward can be obtained only from the controller.

The "problems" to which you refer are already with us. Few people know how to deal with others on any basis other than power and control. In part, this ignorance comes about through our ability to use language to disguise nastiness as niceness. We are currently being exhorted by some politicians to put an end to the welfare system (giving people what they need unconditionally) and "helping" them become self-reliant (begin earning their own money under socially established conditions). If they are unable to reorganize so as to become skillful workers before they starve to death, too bad. After they are dead, only those capable of the required reorganization will be left, which will prove that the politicians (and Darwin) were right.

There's a certain amount of primitive sense in this approach. Clearly, one wants to be among those who have the power to establish contingencies, not those subject to them. Since everyone sees the same advantages in gaining power, it's necessary for a certain amount of ruthlessness to prevail, backed up by a willingness to engage in whatever violence is necessary. The society is defined in terms of winners and losers, so it's better if one becomes a winner. If you don't learn to compete, you're a lamb for the slaughter; sooner or later someone else will be deciding whether or not you eat.

It's no wonder to me that we live in the world's most violent society.

Wayne Hershberger (920206.1400) --

>The problem to which I think Joel Judd was alluding concerns the

>interface between the whole model and its metaphysical context (boss
>reality). Where does one draw THAT dotted line?

Between the model and that to which it alludes (i.e., it is drawn in the whole plane of the paper).

I don't see the problem here; never have. I steadfastly refuse to believe that you're claiming that the totality of What Is is contained in our models or even in our perceptions. To me, it's commonplace to realize that I don't understand or even see all of what is going on. It doesn't seem strange to assume that some aspects of what is going on are unavailable to human perception. I learned this not from philosophy but from experience. I've tried repeatedly to draw analogies that show how matter-of-factly we take for granted the existence of order and interaction beyond the boundaries of what we can experience -- the simplest was flipping a switch on the wall and seeing a light go on in the ceiling. But you've never seemed to understand what I was getting at, perhaps because you're making it more complex than I intend it to be. I'm saying nothing more than that under the surface of a table, we assume there is more material of the same kind, but don't experience it. I'm saying that if we pull the plug in the bathtub, we assume that the water is going somewhere and not just disappearing, but we don't see that place.

The world is full of apparent causes and effects with no indication of how the cause influences the effect. To me, science consists of assuming that there is a hidden reality in which the explanation can be found, and then making and testing systematic guesses about what the link could be -- making models. Occasionally we find ways to uncover what is happening behind the scenes, and check our guesses by direct comparison with experience. Usually we then find that our models weren't quite right, but weren't totally off the mark, either. And of course we also find that we have merely pushed the barrier back a little bit: beyond what we observe there are still hidden links.

The boss reality, which I repeat is hypothetical, includes all that we experience and all that we do not experience. We understand some of it, as it is represented in human perception, but not all of it, and certainly not any of it as it might be represented in some other perceptual system that interacts with the Immanent Order in a different way.

Rick Marken (920206.1423) --

>Won't anyone tell me what "cognitive penetrability" is? Or do I have
>to go to the library?

It means the property of being penetrable by cognition. If something is impenetrable to cognition, then you don't have to blame yourself for not understanding it; it's nature's fault. It also means that you're free to make any guesses you like about the impenetrable thing, because nobody will ever be able to check up on you, at least not cognitively.

If you go to the library and look up Pylyshyn, you won't learn a lot more than that.

It seems to be good for you to be driven relatively nuts once in a while (a short drive). You put your finger on what's really wrong with the Beer cockroach approach:

>I think that the discussions of the various clever mechanisms that might
>be used to produce interesting "behavior" are obscuring a basic question
>-- is behavior a process of controlling variables??

It's possible that by accident, Beer has found some neural circuits that actually control the behavior of a cockroach in some respects. I say "by accident" because if you define behavior in terms of actions and not purposes, it's a matter of luck whether you pick an action or a measure of action that happens to be closely associated with the intended result. Look how long neurophysiologists have known about the stretch and tendon reflexes without being able to figure out how they enter into ordinary behavior. Knowing the circuitry doesn't help much if you don't know what it's for. To judge what it's for for a cockroach in terms of human perceptions from the laboratory frame of reference is to anthropomorphize inappropriately.

I'm getting the Beer book, and when it's ready, Pat and Greg's version 4. Why don't you do the same? I think we'll both be able to evaluate the Beer model better when we've played with it a while from the CT point of view.

Oded Maler (920206.1934) --

More on Beer model:

>Too much, to my taste, is devoted to the simulation of the neural nets,
>and too little is devoted to the dynamics of the interaction with the
>physical environment, which is done mostly on qualitative-level.

I agree. The relationship "velocity = k*force" is completely ad-hoc, and doesn't reflect any real physical dependency. At the very least, if the claim is that velocity is limited by air viscosity, the velocity should enter as the square. And I don't see any way to introduce the effect of a force applied externally to a limb: that would involve the muscle spring constant and a model of the muscle. You'd need to know how the leg would move under an external force with constant neural input.

One interesting aspect of the model (as much as I remember of it) is the interconnection between leg-activating neurons. This has a resemblance to what is found in snakes, worms, fishes, and centipedes. A signal at the start of the chain creates a wave that travels to successive segments or legs, which move in a sort of sine-wave with phase delays as you go down the chain. I could accept this as a very complex output function involving many muscles, with cross-connections that create forward locomotion (and maybe backward too) without the need for detailed feedback. If that's the case, we wouldn't expect this locomotive function to be capable of any variations except in speed of operation. An organism with a flexible body would turn just by bending the body, not by changing the relative speeds of left and right chains. In the cockroach, however, turning can be accomplishes only by varying the speeds on left and right, so we'd expect the left and right chains to be separately adjustable for speed, and perhaps amount of movement per cycle. This "output function" could be affected in several different ways by signals from higher-order systems.

Pat and Greg Williams (920207.0528) --

>One of the "modes" of behavior of Beer's bug is "edge-following." What
>happens is when an antenna hits a wall, it CAUSES the bug to turn away
>from the wall and also activates a slower- acting circuit which turns
>the bug back into the wall (meanwhile having moved forward), and the
>process repeats. Thus, the bug "bounces" its way along a straight edge,
>via a process incorporating both locally open-loop (no adjustment with
>respect to perceptual changes) AND closed-loop (the "kick" away from the
>wall depends on the antenna signal) processes which are happening in
>part simultaneously.

I'm beginning to itch to get my hands on this model of yours. Maybe you could try some things before I do.

One change I would make would be to make the antenna signal proportional to amount of antenna deflection. Then each antenna signal would increase the speed (or stride) of the legs on the same side as the antenna, and decrease the speed (or stride) on the other side, causing the path to bend and reduce the deflection of the antenna. So the control system would be controlling the signal indicating antenna deflection, keeping it near zero, as the bug moved forward. Now you don't need the separate circuit to move the bug back toward the wall, and it won't "bounce" along the wall (if the gains are set appropriately) but will follow it smoothly.

This would also work with only an on-off antenna signal, but control of the signal would be much cruder. You still wouldn't need a special circuit to turn the bug back into the wall -- seeking the food would take care of that. If the bug weren't trying to get to food on the other side of the obstacle, why would it ever turn back toward the wall, anyway? It doesn't care about "edge-following" per se, does it? Actually, to make the bug follow an edge all you would have to do would be to set the "touch" reference signal a little higher than zero.

In any case, this is not an open-loop cause-effect system: bending of the path controls the antenna signal. It's a control system. When you look at only the input-output part of a control loop, of course you see input causing output. But the real causation runs the other way: output controls input.

I presume the smell-detectors are arranged so that total smell error (sum of smell signals from left and right antennas subtracted from a reference-amount indicating hunger) drives the overall speed (left and right legs) and the difference in smells increases and decreases the stride of the left and right legs oppositely to bend the path. The control systems try to keep the left and right smell signals equal and to bring the sum of the signals to some large amount. That's about the minimum design, and I presume cockroaches aren't any more complex than necessary.

So we have two more control systems, controlling the sum and difference of amounts of smell from the left and right antenna smell receptors. Still nothing open-loop.

If I were designing this cockroach's motive system for maximum generality of use, I'd build a steering circuit that took care of lengthening the stride on one side and doing the opposite on the other side, and a velocity system for speeding or slowing both sides simultaneously. In such a small animal, control of these variables might not be necessary -- these would just be two output functions. Then the signals entering these two output functions could come from multiple sources -- control systems controlling for touch, for smell, for light, and for other things. You'd then have almost a direct analog of the control systems in the "Gatherings" (formerly "Crowd") program. In fact, there's no reason but computer power that you couldn't incorporate the bug control systems into the Gatherings program and have a kitchen full of cockroaches avoiding each other and obstacles while heading toward their respective goals (food for hungry ones, water for thirsty ones, pheremones for horny ones, etc.?).

Tell Pat to hurry up with Version 4.

Best to all,

Bill P.

Date: Sat Feb 08, 1992 12:47 pm PST
Subject: credit and control

[from Joel Judd]

Wayne and Bill (920207/8)

In the foreign student writing courses we always ahave a little unit about plagiarism, so I would be remiss if I didn't point out that the model interface comments Wayne mentioned and Bill repeated stemmed from Mark Olsen. There, now my conscious is clear. And since we're here...

Bill (920208)

You would be a big hit at a sociology of education convention. Or any education convention for that matter.

Re: Education--Isn't "defining the result" an ambiguous statement from a PCT point of view? I assume that you are not referring to particular behaviors, but to socially accepted goals or attitudes reflected in particular behavior. Otherwise, we're talking about the traditional "course objectives," "class objectives," etc. that every good teacher is taught to determine ahead of time and which are typically expressed in terms of student behavior.

As I finish the last chapter of the dissertation, I'm faced with proposing some pretty strange ideas; namely, that teaching something like language means helping students develop L2-like perceptual references. This is no longer language per se--it's a way of interpreting the world. On top of that, it's not a clear alternative to the L1, it is mixed up with L1 perceptions. In addition, there is no guarantee one will learn. This is also a problem with Education in general. What do you do with the people who simply DO NOT LEARN? When do we admit that we (if the shoe fits...) cannot promise learning? Finally, maybe those definitions of learning most tied in with critical high level social goals promote longer lasting reorganization; that is, reorganization is more persistent. I'm thinking of the international TOEFL cartel, which through force of history and use has become THE definition of English learning. People who otherwise have abysmal (in my subjective judgment) language skills manage to keep going until they get a sufficient score on the TOEFL to get into a US university or whatever they're using it for. I'm grossly overgeneralizing, of course, but--good or bad--it just seems to constrain the environment more than other language goals.

Re: Welfare--Do I understand that of the two alternatives mentioned, you prefer the unconditional option?

Date: Sat Feb 08, 1992 7:36 pm PST
Subject: PCT BUG: GO FOR IT!!!

From Pat & Greg Williams

>It's possible that by accident, Beer has found some neural circuits that
>actually control the behavior of a cockroach in some respects. I say "by
>accident" because if you define behavior in terms of actions and not
>purposes, it's a matter of luck whether you pick an action or a measure
>of action that happens to be closely associated with the intended result.

We don't think Beer ONLY defines behavior in terms of actions, although he doesn't have a PC-theoretical underpinning for his modeling. It appears that he set out to model some "behaviors" (his word) like "satisfying hunger," which is a PCT-type "true" behavior (and which, in Beer's model, can indeed occur via varying-in-the-face-of-disturbance, PCT-type "actions." But he isn't consistent and thorough in keeping "actions" and "behaviors" separate. So, we say (as we said to Rick) ENLIGHTEN HIM and SHOW HIM THE ADVANTAGES of a PCT foundation for his future modeling!

There is a long history of PCTers pointing out the errors of others' behavioral models. That practice certainly hasn't resulted in widespread appreciation of PCT. Perhaps it is time to talk about the relation of PCT to what is CORRECT (i.e., it controls) about a model which has actually generated considerable interest among non-PCTers BECAUSE of its PCT-correctness (a fact which needs to be explained by PCTers to the non-PCTers showing interest).

>One interesting aspect of the model (as much as I remember of it) is the
>interconnection between leg-activating neurons. This has a resemblance to
>what is found in snakes, worms, fishes, and centipedes. A signal at the
>start of the chain creates a wave that travels to successive segments or
>legs, which move in a sort of sine-wave with phase delays as you go down
>the chain. I could accept this as a very complex output function
>involving many muscles, with cross-connections that create forward
>locomotion (and maybe backward too) without the need for detailed
>feedback. If that's the case, we wouldn't expect this locomotive function
>to be capable of any variations except in speed of operation. An organism
>with a flexible body would turn just by bending the body, not by changing
>the relative speeds of left and right chains. In the cockroach, however,
>turning can be accomplishes only by varying the speeds on left and right,
>so we'd expect the left and right chains to be separately adjustable for
>speed, and perhaps amount of movement per cycle. This "output function"
>could be affected in several different ways by signals from higher-order
>systems.

A while back on the net there were pleas for better empirical grounding of discussions. We advise you to read Beer's book to discover how thoroughly wrong your memories are. The gait circuitry of Beer's bug is completely different than you suppose. And for our next homework assignment...

>I'm beginning to itch to get my hands on this model of yours.

GO FOR IT!!!!!! (See below at additional exclamation points.) We've wanted to do a PCT-based bug for months, but haven't had time. PLEASE read Beer's book first, though.

>Maybe you could try some things before I do.

Not likely... we've got two books to publish, remember?

>If the bug weren't trying to get to food on the other side of the obstacle,
>why would it ever turn back toward the wall, anyway? It doesn't care about
>"edge-following" per se, does it?

(Exasperated.) READ THE BOOK! (Hint: the answer to the last question is yes for real live cockroaches.)

>In any case, this is not an open-loop cause-effect system: bending of the
>path controls the antenna signal. It's a control system. When you look at
>only the input-output part of a control loop, of course you see input
>causing output. But the real causation runs the other way: output
>controls input.

Glory be, Beer must have "guessed" right. We think this is reason for jubilation, not commiseration! HE'S ON YOUR SIDE. If you still think he's a little confused, help to set him straight. We think you'll be doing YOURSELF a big favor.

>Tell Pat to hurry up with Version 4.

It isn't to the Pat (programming) stage, yet. Greg is still scouring the literature for information on sensory transducers, interneuron input-output properties, known control circuits, etc., etc., for a plethora of invertebrates. Hooray for Bullock and Horridge and other compendia!

!!!!!! But, lucky you, you don't really NEED Version 4, since you know how to program in C, and Version 3 comes complete with Turbo C source code. !!!!!!

You have a copy of Version 3, don't you? We can USPS it if you don't. Pat is available to consult on the obscure parts. Better do it to NS87, the floating-point version of the program, since it doesn't have all the scaling crap.

(Of related interest: We just sent back Borland C++ (for DOS and Windows), which is too damn big for Pat's 386SX with 1 meg of extended memory and a 40 meg hard disk. If we installed the whole thing, it would take over 40 megs. And we couldn't compile some modules of PictureThis because the compiler didn't get out of the way -- it is over 1meg itself, and you'd think it would swap part of itself out, but the book says "you need more extended memory." Nope. We're sticking with Turbo C (actually Turbo C++, but we don't use objects). Remember that every IBM clone with Windows can still run DOS programs (if the operator isn't DOS-prompt-shy), but nowhere near every IBM clone can run Windows programs! We're JUST SAYING NO to Windows programming.)

Pat & Greg

Date: Sun Feb 09, 1992 8:47 am PST
Subject: Education; Welfare; BEERBUG

[From Bill Powers (920209.0900)]

Joel Judd (920208.1321) --

Thanks for correction and apologies to Mark Olsen.

>Re: Education--Isn't "defining the result" an ambiguous statement from a
>PCT point of view?

Yes. I was thinking of it from the normal point of view. A result is a test with answers filled in. While it is true that many teachers are able to present materials in a way that makes learning easier and quicker, there isn't any method that I know about to make students eager to learn or to get them to value learning. The main incentive, as I've experienced and seen (a limited sample), is fear of failure or of losing out in competition. Course objectives concern what the student has to do in order to avoid failure or graduate with impressive grades. When meeting the course objectives becomes only a means toward controlling for artificial goals, as soon as the requirement for meeting them in that way is removed the course objectives are forgotten. If the course objectives were stated in terms of the students' real goals, they would remain relevant.

>Finally, maybe those definitions of learning most tied in with critical
>high level social goals promote longer lasting reorganization; that is,
>reorganization is more persistent.

I get the idea but it could be put less misleadingly. The outcome of any series of reorganizations is simply the last organization that existed when reorganization ceased. A "longer lasting reorganization" is not longer-lasting because reorganization was extra-effective. It lasts longer because it succeeds in preventing further intrinsic error. The reorganizing system doesn't know the difference between a "final" organization that lasts 10 minutes and another that lasts 10 years.

I'd say it this way: definitions of learning most tied in with the goals people actually have (as opposed to those forced on people as a means of reaching their actual goals) are most likely to promote satisfaction of intrinsic goals, and thus reduce the need for further trial-and-error learning.

>Re: Welfare--Do I understand that of the two alternatives mentioned, you
>prefer the unconditional option?

Certainly. In America, there shouldn't be any particular prestige factor in getting enough to eat, staying warm, and having shelter from the elements. We don't punish criminals by starving them, freezing them, or making them live outdoors in snow and rain. We punish only poor people that way, for the crime of having no skills, no ambition, no confidence, and no hope. The welfare proposals now being offered by the less sympathetic simply require that poor people get jobs, leaving it up to their skills, ambition, confidence, and hopefulness to find a way to do so -- at a time when there are hundreds of applicants for every job offered. I think that there should be a floor under poverty at the highest standard of living that is feasible, without any qualifying requirement other than breathing. If business can't manage its affairs well enough to provide even elementary survival for everyone in this country, then business is incompetent and we need a different system.

When you leave everything up to poor people's reorganizing systems, you shouldn't be surprised if some of the solutions they discover are socially unacceptable. Reorganization ceases when an organization is found that works to reduce intrinsic error. If that entails holding up 7-11s and selling dope, the reorganizing system doesn't care.

I don't give a damn about economic arguments. There are enough skilled people in this country, and there is enough industrial savvy, to provide a comfortable life (if not necessarily interesting) for every American, and then some. All that's needed is the idea that doing so might be satisfying and productive of future gains. Or just the right thing to do.

Greg & Pat Williams (920208.1955) --

I'm willing to say that Beer is investigating control phenomena. I'm willing to say that B. F. Skinner was investigating control phenomena. In fact, one of the things I've been saying is that EVERYONE has been investigating control phenomena. I also agree that it would be nice to convince everyone that using control theory is the best way to do this. If Beer is amenable to adopting control theory, great. Is he? You're in touch with him, aren't you?

>There is a long history of PCTers pointing out the errors of others'
>behavioral models. That practice certainly hasn't resulted in widespread
>appreciation of PCT.

You're right. I wish others had more appreciation of seeing mistakes pointed out -- I wish scientists really behaved like scientists. On this net we don't seem to have a high level of defensiveness. When errors are pointed out (validly) the reaction is usually "Oh, you're right, thanks." I haven't seen much of this in "normal" science. That's why I like communicating through CSGnet with people who voluntarily signed on to control theory. Even when they tell me I'm wrong.

>>But the real causation runs the other way: output controls input.

>Glory be, Beer must have "guessed" right. We think this is reason for
>jubilation, not commiseration! HE'S ON YOUR SIDE.

When he says so I'll believe it.

Re: C++

I can't even run Windows with my little 20 MB disk, which I have to keep cleaning out. I dropped Magellan to make room for PictureThis V.4. I'm using C version 1.5, which seems to do everything. I haven't got into objects yet and probably won't unless someone can tell me what they do that good programming doesn't do. I don't much like bundling data sets with operations on the data. Seems limiting to me. Also, the implication is that objects contain the operations done to them, which doesn't make much sense to me, modelingwise.

By the way, PTR4 is a great program, folks, which can now do full typesetting -- the 180+ page manual for the registered version was done with PictureThis. Only one lack: a crib sheet organized by function so you have some idea of what to do next when faced by a blank screen with an X in the middle of it or when you're partway through something and forget what the next step is. I prefer the uncluttered screen devoted completely to the drawing, so this isn't a complaint. Learning goes very fast because Pat has made the commands simple and logical.

In case my private post didn't make it, I do need a copy of version 3 of the neural net program. The old one is probably here somewhere, but ...

Best

Bill P.

Date: Sun Feb 09, 1992 4:41 pm PST
Subject: Emotions again

[from Kent McClelland (920209)]

Bill Powers (920203.1000)

Sorry to take a while to get back to you, but I'm only now catching up (more or less) on last week's net postings.

Thank you for your words on the subject of emotion. I shared the post with my seminar students, and they thought it cleared up a good many things, but, as is typical, not everyone was satisfied. In particular, one student reacted to these words. . .

> Many emotions can be
>interpreted by asking not what it feels like but what you want to do when
>you feel it.

by saying that the real issues is where the "wants" come from. Another student averred that emotions have a more central role in guiding behavior than you seem to allow. They are, she suggested, like a "12th level" of the hierarchy that provides reference signals for the levels below.

While I didn't see how the "12th level" suggestion could make any sense within the HCT framework, it did set me to thinking when I went back and read your original post juxtaposed with a later post on "intrinsic error." (Bill Powers 920203.1900). Maybe we might consider emotions as part of the intrinsic error regulation system. Emotions might be perceptions that bring news of intrinsic error or its resolution by successful reorganization. Newborn babies certainly seem to come equipped with emotions, especially the negative ones like rage, fear, and general discomfort. Interestingly, it apparently takes them awhile (like the four to six weeks before smiles appear) to begin to register positive emotions. Emotions also often seem to be accompanied by a greatly increased rate of reorganization, particularly of the flailing around variety, and pain can start us on reorganization in a hurry.

I have some other thoughts on this but no time to send them tonight.

Mary Powers (920203)

>Bill has had two checks from Grinnell dated a week apart

That's interesting. Trust our Treasurers office to combine inefficiency with its penchant for red tape. I'll make an inquiry (and in the same time see if they can get my wife's annuity withholding straightened out!) I think you were probably only supposed to get one, because we've just installed and used the programs this semester. However, if I were you I guess I'd take the money and run.

Be well,

Kent

Kent McClelland   Office: 515-269-3134
Assoc. Prof. of Sociology   Home: 515-236-7002
Grinnell College   Bitnet: mcclel@grin1
Grinnell, IA 50112-0810

Date: Sun Feb 09, 1992 7:00 pm PST
Subject: The Big Picture

[from Gary A. Cziko 920209.2040]

To Bill Powers, Pat & Greg Williams, Rick Marken, Wayne Hersheberger, Hugh Petrie, Chuck Tucker, Clark McPhail and any other PCTers interested in the Big Picture and computer simulations:

Ever since I first got interested in psychology I've been trying to develop a general understanding of what behavior is all about and how it changes and develops over time.

Stage 1: At first as an undergraduate I was introduced to the standard S-R perspective and was much attracted to Skinner's ideas. I remember watching the fish in my aquarium and thinking to myself, "If we could just understand all the input variables and the organism's history, we could make accurate predictions of EACH movement that that little guppy makes."

Stage 2: As I became more aware of the difficulties of predicting and controlling behavior, I settled for a statistical perspective. There are just too many variables interacting in too many complex ways for us to make good predictions, so we have to be happy with accounting for percentages of group variance. Better and better instruments coupled with more and more sophisticated computer analyses should, however, permit better and better predictions, understandings, and theories, or at least it seemed to me at the time.

Stage 3: I then discovered quantum mechanics and chaos theory and everything went to hell. If we couldn't make accurate predictions for things as simple as electrons and dripping faucetts, then how in the world could we make predictions for humans and other animals?

Stage 4: Perceptual Control Theory then came along and things started to fall in place. People do not normally act chaotically. The variation in their behavior are simply ways of cancelling out disturbances which would otherwise prevent them from reaching and maintaining their internally specified goals. Organisms are "chaos neutralizers." If we knew exactly what a living control system's goals were and what disturbances it would meet, we COULD in principle make valid predictions concerning behavior. Until, that is, long-term error was encountered and reorganization started taking place. Now we're back into the chaotic, messy, evolutionary, unpredictable part. If there is only one way to reorganize to reduce intrinsic error, then we could probably still predict where the reorganization would lead too, although we still couldn't predict how long it would take for an organism to get there (although statistics might be of some use here). But if there were many ways to reorganize to reduce the intrinsic error, then it would be difficult if not impossible to predict how reorganization would take place and subsequent behavior.

I have two questions about all this. First, how is my Stage 4 perspective consistent or inconsistent with others who are trying to use PCT for developing a Big Picture?

Second, would it be possible and useful to create a computer simulation that would demonstrate the principles of both control and reorganization? I have only very vague ideas of how this would be done now, but perhaps it could be started by modifying the Gather (nee Crowd) program of Powers, McPhail and Tucker.

The units (people) in the Gather program can now only control; they cannot reorganize. But would it not be possible to allow them to reorganize in some simple (but blind) way when their error is not reduced beyond a certain level (the error could be introduced by changing the environment in some way)? And wouldn't it also be possible to vary the degree of convergence that reorganization would likely lead to? In a "convergent" environment (like first language learning; rats in a maze?), all control systems would eventually come up with the same type of reorganization, but in a "divergent" environment (like second language learning; people in a modern city?) many different reorganizations would be expected either by the same simulated organisms in successive runs or by a number of similar organisms in the same run.

I realize that the demos that Bill has created up to now are attempts to model actual behavior (as the control of perception, of course) and so what I am considering here is a bit different. But it may nonetheless be useful to give a type of Big Picture as to what PCT means for developing organisms and their behavior. An interesting interplay of "in principle" predictable behavior whose sole purpose is to control variables interspersed with reorganizations of variable predictability according to the "convergence" or "divergence" of the situation. At least this is the Big Picture that I seem to be developing and would like to get reactions from others to how their Big Picture overlaps (or doesn't) with mine.--Gary

P.S. In case anyone would want to contact Randy Beer and his bug via e-mail, his address is beer@cthulhu.ces.cwru.edu (Greg, this is different from the address I gave you earlier, although the earlier address appears to work as well).

Gary A. Cziko   Telephone: (217) 333-4382

Date: Mon Feb 10, 1992 1:46 am PST
Subject: On bugs, psychologists and roboticists.

[From Oded Maler 92.02.10]

Rick's problem can be summarized as:

1) Most of the people who are interested in *expalining* behavior of living systems (psychologists), think in categories that are too far from your "conrol" point of view. They are interested in the *same question* as you are but approach it using totally inadequate tools.

2) Some people who are trying to *build* working artificial systems (robotoicists), either real (Brooks) or simulated (Beer) are interested in a *different question*, which is related to the previous one, because apparently the only systems that currently work are living ones. Their aproach is much more close to the "control" approach because they are forced to think in terms of sensors and actuators. For them, if they succeed in making something that works, the question what does it "really" control is not relevant. You can poke elecetrodes into their (creatures') sensors and tell them "your creatures really controls for this and that perceptual variable", maybe you can do it even directly by looking into their circuits and programs. So what? You can sell them your ideas only if you show them that by using the PCT ideas you can design systems that achieve complex behaviors which are otherwise unachievable. This is a question of design methodology, not of a scientific theory.

If you keep in mind the distinction between the two questions, I think you'll get more positive feed-back (if this is still your controlled varible (?)) from the second group than you seem to be getting from the first.

I'd appreciate:

1) A copy of the simulated insect program (Gregg&Pat).
2) A free copy of Closed Loop (if it's not to much to ask).
3) A list of regular mail addresses of CSG-members (Gary?)

I'm also glad to report that a copy of Living Control Systems has finally arrived to this part of the world, and if I will not be enlightend within the next few months, I will be the only one to be blamed.

Best regads

Oded Maler
IRISA
Campus de Beaulieu
35042 Rennes
France

Date: Mon Feb 10, 1992 4:04 am PST
Subject: Going for it

From Pat & Greg Williams (920210)

>Rick Marken (920209)

>I will take it as an action to read his book more thoroughly and write a
>report on it from a behavioral science perspective.

Wonderful! We are hoping that your report will concentrate on two main points:

(1) Beer's bug as a working simulation embodying the principles of hierarchical control of perceptions (it is robust BECAUSE it embodies those principles); (2) such simulations can be improved and more easily designed by adopting a consistent underlying theory based on hierarchical control-of-perception ideas. And we hope that you will emphasize convergences of PCT ideas and Beer's work, rather than divergences.

Rick, as you get started, also take a look at the Beer, et al. article in AMERICAN SCIENTIST, Sept.-Oct. 1991. It is the only place you'll find a nice drawing of the overall bug nervous system circuitry.

>Oded Maler 92.02.10

>You can sell them your ideas only if you show them that by using the PCT
>ideas you can design systems that achieve complex behaviors which are
>otherwise unachievable.

Agreed!

>If you keep in mind the distinction between the two questions, I think you'll
>get more positive feed-back (if this is still your controlled varible (?))
>from the second group [roboticists] than you seem to be getting from the
>first [psychologists].

We think so, also.

>I'd appreciate: A copy of the simulated insect program (Gregg&Pat).

The bug is in the (air) mail.

Best,

Pat & Greg

Date: Mon Feb 10, 1992 4:46 am PST
Subject: SAB92 Announcement

Conference Announcement and Call For Papers

FROM ANIMALS TO ANIMATS

Second International Conference on Simulation of Adaptive Behavior (SAB92)

Ilikai Hotel

Honolulu, Hawaii, December 7-11, 1992

This conference is the successor to SAB90 - which was held in Paris in September, 1990. Its object is to bring together researchers in ethology, psychology, ecology, cybernetics, artificial intelligence, robotics, and related fields so as to further our understanding of the behaviors and underlying mechanisms that allow animals and, potentially, robots to adapt and survive in uncertain environments.

The conference will focus particularly on simulation models in order to help characterize and compare various organizational principles or architectures capable of inducing adaptive behavior in real or artificial animals.

Contributions treating any of the following topics from the perspective of adaptive behavior will receive special emphasis.

Individual and collective behavior	Autonomous robots
Neural correlates of behavior	Hierarchical and parallel organizations
Perception and motor control	Emergent structures and behaviors
Motivation and emotion	Problem solving and planning
Action selection and behavioral sequences	Goal directed behavior
	Neural networks and classifier systems
Ontogeny, learning and evolution	Characterization of environments
Internal world models and cognitive processes	Applied adaptive behavior

Submission Instructions

Authors are requested to send two copies (hard copy only) of a full paper to each of the Conference co-chairs (Meyer, Roitblat, & Wilson). Papers should not exceed 10 pages (excluding the title page), with 1 inch margins all around, and no smaller than 10 pt (12 pitch) type (Times Roman preferred). Each paper must include a title page containing the following: (1) Full names, postal addresses, phone numbers, email addresses (if available), and fax numbers for each author, (2) A 100-200 word abstract, (3) The topic area(s) in which the paper could be reviewed (see list above). Camera ready versions of the papers will be required after acceptance. Computer, video, and robotic demonstrations are also invited. Please contact Herbert Roitblat to make arrangements for demonstrations. Other program proposals will also be considered.

Conference committee

Conference Chair

Jean-Arcady MEYER
Groupe de Bioinformatique
URA686.Ecole Normale Superieure
46 rue d'Ulm
75230 Paris Cedex 05
France
e-mail: meyer@wotan.ens.fr
meyer@frulm63.bitnet

Herbert ROITBLAT
Department of Psychology
University of Hawaii at Manoa
2430 Campus Road
Honolulu, HI 96822
USA
email: roitblat@uhunix.bitnet,
roitblat@uhunix.uhcc.hawaii.edu

Stewart WILSON
The Rowland Institute for Science
100 Cambridge Parkway
Cambridge, MA 02142
USA
e-mail: wilson@smith.rowland.org

Organizing Committee   S. Gagnon, H. Harley, D. Helweg, M. Hoffhines,

Program Committee

A. Berthoz, France	M. Bitterman, USA
L. Booker, USA	R. Brooks, USA
P. Colgan, Canada	J. Delius, Germany
S. Goss, Belgium	L. Steels, Belgium
R. Sutton, USA	F. Toates, UK
S. Tsuji, Japan	W. Uttal, USA
D. Waltz, USA	

Official Language: English

Important Dates

JUL 15, 1992	Submissions must be received by the organizers
SEP 1, 1992	Deadline for early registration
OCT 1, 1992	Notification of acceptance or rejection
NOV 7, 1992	Deadline for regular registration
NOV 15, 1992	Camera ready revised versions due
DEC 7-11, 1992	Conference dates

Date: Mon Feb 10, 1992 8:53 am PST
Subject: bulletin board

I've heard there is a control systems group email
bulletin board. I'd like to be added to the list.

Thanks,

Dan Stone


```
----------------------------------------+
HIGH TECH  |   DSTONE@vmd.cso.uiuc.edu   |
-----------+----------------------------|
MED  TECH  |   (WORK) 217-3334537    |
           |   (HOME) 217-3285957    |
-----------+-------------------------
LOW  TECH  |   DAN N. STONE          |
           |   UNIV. OF ILLINOIS     |
           |   DEPT. OF ACCOUNTANCY  |
           |   AND INFO. SYSTEMS     |
           |   BOX 12                |
           |   1206 S. SIXTH ST.     |
           |   CHAMPAIGN, IL 61820   |
```

-----------+-------------------------+

Date: Mon Feb 10, 1992 8:54 am PST
Subject: Welcome to CSGnet

Welcome to the Control Systems Group Network!

As a new subscriber, you will be joining a discussion already well underway. To help you to get your bearings, I have attached here two documents. The first is a short introduction to control theory by Bill Powers. The second is information on books relevant to control theory. Bill Powers has also developed some very useful computer demonstrations of control theory which run on any IBM-compatible PC with a mouse. For information about these you should contact him directly at POWERSD@TRAMP.COLORADO.EDU

Please do not hesitate to introduce yourself and your interests to the network. And if you don't find the current discussion relevant to your interests, just let us know what they are; if they have anything to do with the "life sciences" broadly conceived, you can almost be guaranteed to get some interesting responses from our subscribers. I strongly recommend this action instead of just signing off before learning what a control theory perspective of your interests would be like.

Finally, please remember that commands should be sent to the listserver (LISTSERV@UIUCVMD or LISTSERV@VMD.CSO.UIUC.EDU) and NOT to CSG-L.

Again, welcome to CSGnet and the fascinating world of (perceptual/hierarchical) control theory.--Gary

Date: Mon Feb 10, 1992 8:58 am PST
Subject: one last regorganization [for now]

[from Joel Judd]

Bill (920209)

>A "longer-lasting reorganization " is not longer lasting because reorganization
 was >extra-effective.

Sorry, I didn't mean to imply a "Contac" (tiny time capsules) view of reorganization. What I had in mind dealt with the fact that many people simply learn better that others. So what I was trying to say/ask had to do with those who, by their nature and/or effort, seem to deal with something like language learning more" completely" and effectively. For them, there are higher level goals involved (communication, understanding, friendliness, etc.) and satisfying these may take awhile and require everything from almost native-like pronunciation to proper levels of address to appropriate non-verbal gestures.

Now I remember recently saying something to the effect that higher-level goals could be satisfied WITHOUT such things as good pronunciation. Hmmm. There seem to be two types of characterizations here. There are those who reorganize because there is intrinsic error arising from their attempts to reach their goals, and there is intrinsic error because others are forcing these to deal with goals not their own. The former are more likely to effectively reorganize?

Looking at the last part of Runkel's book relevant to research helped a little bit with this problem. In accounting for variability among peoples' actions, there are three sources to keep in mind. I assume these apply to reorganization as well since the reorg system has to work through the hierarchy. FIRST, environmental factors must be considered as disturbances. In the case of reorg, disturbance must persist until chronic. If the extant hierarchy does something to acceptably reduce error, reorg (hence "learning") will not come into play. SECOND, variability will depend upon the "means we choose to oppose the disturbance." These will depend upon three related factors: (a) what solution(s) we think will be successful, (b) whether such solution(s) implies conflict with another goal, and (c) what solutions we can conceive of in the first place. This second source of variation seems to be the most complex, especially when trying to understand another, because it deals with past learning (memory) and our previous experience with similar perceptions, and our general problem solving skills, expertness, imagination, or whatever you want to call it. THIRD, the environmental allowances for particular courses of action (e.g. you're not gonna swim your way out of the Sahara).

Thus, limited experience and perceived environmental options can limit the possibilities for solutions resulting from reorganization-- hence your description of alternatives among the poor. Yet it is interesting how we hold up those who "reorganize" out of a vicious circle (the "Abraham Lincoln Syndrome"?) and "make it," as the way everyone should be able to do it.

Date: Mon Feb 10, 1992 8:59 am PST
Subject: through a glass, darkly

[From: Bruce Nevin (Mon 920110 08:44:46)]

Bill, re semantics:

Now we are agreed that there appear to be some control systems (ECSs) specialized for language. How far up or down the hierarchy this specialization extends is an open question. My guess is, not below phonemic segments (contrasts) and not above program, but I am by no means certain of that, and in any case that is a vast terrain.

I imagine this not as two hierarchies but as two lineages or threads of control (for given perceptual input or error/effort output) within a single hierarchy. Maybe a vacuous distinction, but I prefer it as a way of leaving the question open.

Language appears to shadow perceptual reality, to mirror it not perfectly, but by murky approximation, like that of the Hellenistic-era "glass" in which one could see things only "darkly." But the disadvantages due to inaccuracy and loss of detail are balanced by the advantages of a well-defined structure that is a matter of social agreement. On that structure (and by means of it) we mold our agreements about the perceptions to which language refers. Harris says the function of language is error-free transmission of information. I would say rather that this scaffolding function is a more essential characteristic of language, and includes the notion of linguistic information within it.

I believe that your proposal for category perception will serve for the word-level linking of the two control lineages within the hierarchy. I will call this the rebus-linking function of category perception, in honor of your example of |red|, "red," |square|, and "square". Reading of rebuses (rebi?) could be accounted for by some other mechanism linking words with meanings, and perhaps more plausibly: the delay we experience puzzling it out could well be due to translating from the picture of a square to the word "square." Imagine the following vertical column of elements in the upper left-hand corner of an envelope, from top to bottom: a sketch of a stack of cordwood, the name "John," the abbreviation Mass. My father told me around 1952 that this had served as the return address of a person he knew. We lived in Lowell, a city near Andover. What is John's last name? Of course, this rebus depends upon a pun as well as on the proposed rebus-linking. But is not rebus-linking itself a kind of pun? More extreme-seeming is the linking of a color with the typographical shape of each letter of the Latin alphabet--a form of synaesthesia recently mentioned on the Linguist Digest.

But rebus-linking also buys a simple account of zeroing along lines you have repeatedly advocated. If either |square| (the perception of a square configuration) or "square" (the word) will satisfy the input requirement of a category recognizer, then socalled zeroing is when the nonverbal input |square| satisfies the input requirement, and so the word "square" need not be present. It is not quite so simple--there are some input requirements of programs that control words and not rebus-linking categories--but putting it in those terms suggests the shape of a possible solution.

Harris doesn't talk about the word dependencies before linearization. Linearization must be under program control. In English, put the word that is "about" the other word or words in second position; put the operator after its first argument. In what sense is one perception "about" another one, in the way that an operator is predicated of its argument? A good answer to that might provide a perceptual basis for the operator-argument distinction along the lines you were suggesting, Bill. Lacking that, it may be that the word classes and word dependencies in language are a new invention, and that by virtue of learning their rebus-linking to perceptual reality we live in a more articulate perceptual reality than do animals lacking language. Harris suggests that the "real-world" dependencies are there, but except in the sense of the arising of an ur-language from real-world (i.e. perceptual) contingencies as sketched in _Language and Information_ I don't see any simple reflection. Any ideas?

Harris doesn't talk about metalanguage statements specifying the reductions. I suppose input signals to program ECSs may come from word recognizers or from category recognizers rebus-linking words to nonverbal perceptions, and this is why it is easier to put programs into words than it is to put other perceptions into words. The inputs to the metalanguage programs for the reductions are from category-recognizers that rebus-link metalanguage words like "word" and "operator" to non-metalanguage words (just as ordinary ones link words to non-word perceptions).

Harris does talk about metalanguage statements of sameness (having same reference) among conditions for reduction. In our PCT perspective, this is not done by metalanguage utterances, but two occurrences of a word having each a rebus-link that threads down the hierarchy from the category level to the same perceptual lower-level signal (real-time or imagined). Can a program determine that the perception rebus-linked with a word has already been rebus-linked with some other prior occurrance of that same word? We need that for the relative clause, for conjunction reduction ("I saw Bill and [I did] not [see] Mary"), and various other reductions.

Can a program determine that an interlocutor is imagined already to know some particular dependency among perceptions, so that the corresponding dependency among words need not be included in an utterance? We need this for common knowledge like using umbrellas to keep off rain and for avoiding repetition in discourse, as well as for reductions that reduce the repetitiousness in sentences. Can you help me to understand how a program-level ECS might get inputs like these?

Lots of thoughts arise from the notion that language is not above experience, but "beside" it, "shadowing" it.

However perceptions are remembered and the memories stored in a living control system, to the extent that word dependencies represent experience, the word-dependencies can also be remembered. As has often been observed, written records constitute an extension of memory, as oral texts do and always have in societies without writing systems for their languages. We use recited formulae to extend short-term memory and as a crutch for long-term memory. Recited texts (written or not) can evoke memories of nonverbal perceptual reality. Furthermore, they can and do also shape the recall of those memories and perhaps the memories themselves. Formulating in language emphasizes some aspects of perception and omits others, and then we fail to remember that it was ever different than the words and associated memory now indicate. I suspect it is in this back-door way that language shapes system concepts, particularly insofar as they are matters of social agreement, rather than explicit self-instruction. But individuals are likely to differ as to the extent of their use of recited formulae to keep their goals in order.

Somewhere in here I think is part of the reason it is so much easier to recognize the application of system concepts to others than to oneself. Gary's recent "counselling" of Rick about applying CT illustrates this. I'm glad you had the time to follow that one up, Gary. But now of course you are a superbly vulnerable target, just as Rick had unwittingly made himself. As an old Indian friend said, when you point your finger at someone else, there are three more fingers pointing back at yourself! CT does change one's perspective. What does CT tell us about why it is so much easier to apply that changed perspective to others than to oneself?

No time to pursue that now, alas.

Bruce
bn@bbn.com

Date: Mon Feb 10, 1992 8:59 am PST
Subject: LOOP II#1 & THE HIERARCHY

SUBJECT: CLOSED LOOP WINTER 1992 & SOCIAL CONTROL

Whatever happened to the Concepts of a Hierarchical Array of Control

Systems? And to the Test of the Controlled Variable?

WHERE is the "Social Control System?" The discussions seem to assume that there is, somewhere, somehow, "A" system that meets the criteria of at least one Control System. A statistical combination of interacting entities is not necessarily "a" control system. There must exist some form of interconnecting structure that serves to define at least one identifiable Controlled Variable, at least one related Output Signal, and at least one related Reference Signal. I think that there are ways to meet these requirements, but I don't see them considered in these papers.

There is discussion of "conflict" from various standpoints, but that only implies interaction between (among?) systems (people) without implying any larger organized form of control system.

To present my current thinking along these lines, I find it necessary to review and refine some of our original work and ideas. That is why my recent offerings may have appeared rather unnecessary and repetitive of earlier discussions.

Let us turn to the simplest description of a control system: The "Logical Box." This is a "Black Box" without the paint.

The Box can take many forms, but careful examination reveals that it has four unique connections between the world outside the box and whatever is within the Box.

Number 1) is normally taken for granted. It is the cord that plugs into the wall outlet. When unplugged, the Box is inert, it responds in no way to any kind of treatment -- or mistreatment. The box requires some source of energy. It may be outside the box or inside the box -- it is essential.

Number 2). This connection, when "disturbed" does nothing. However it is observed that Number 4) changes its condition in a manner related to the disturbance of Number 2). The nature of the relevant disturbance at Number 2) (temperature, pressure, voltage, cosmic-radiation, etc) may be hard to determine. Likewise, the nature of the related changes in number 4) may be hard to identify.

Number 3). This connection similarly affects Number 4) and may respond to a disturbance different from the one affecting Number 2).

Number 4). As long as the Box retains its energy source and BOTH Number 2) and Number 3) remain constant, this connection (Number 4)) is very difficult to disturb. Perhaps impossible short of destruction of the Box.

As described, this box is not a Control System. It does nothing until some specific disturbance is applied to either (or both) Number 2) and/or Number 3). However, if some connection is provided in the ENVIRONMENT of the Box such that the operation of Number 4) serves to reduce the Disturbance of Number 2), then we have a Control System acting to Control (or at least tend to Control) the Disturbance of Number 2). In this case, we find that the same Disturbance applied to Number 3) results in a corresponding change in Number 4) such that the Disturbance of Number 2) is (very nearly) the same as that applied to Number 3). Of course, by suitable modification of the

ENVIRONMENT, the operations of Number 2) and Number 3) could be interchanged or otherwise modified.

This Stylized and Abstract description is consistent with the usual descriptions of Negative Feedback Control Systems. It is presented here to emphasize that such a SYSTEM is DEFINED EXCLUSIVELY in Terms of its INPUT and OUTPUT Characteristics as it interacts with its ENVIRONMENT. In addition, to act as Control System it must respond to another INPUT, serving as a REFERENCE SIGNAL and originating outside the System. Note that this says nothing about the nature of any of these three connections, of the environment, nor of how they may be interconnected inside the Box.

This Box acts ("behaves") by controlling its PERCEPTION which consists of the DISTURBANCE to which it responds and Controls. The TEST OF THE CONTROLLED VARIABLE applies directly to the Box.

As presented here, some important items are capitalized. This will be clarified later by discussion of their relations and implications in terms of the Concept of a Hierarchical Structure of Negative Feedback Control Systems.

Regards, Bob Clark

Date: Mon Feb 10, 1992 9:20 am PST
Subject: Introduction (I hope!!)

Our mailer has been flakey of late so don't know if this will make it, but here goes!

Hi -

As this is my first post, I thought I would introduce myself and my interests to the group.

First let me thank Greg Williams for giving me a pointer to CSGnet. I have been listening in for about a week now and Greg was right it is a VERY active group and I might add VERY interesting.

Second let me be up front about my lack of knowledge about control systems or Dr. Powers adaptation of it to studying behavior. Also, I do not come with any formal training is psychology or linguistics, but with training as a mathematician who has been working in AI for nearly a decade now -- Boy do I feel like a fish out of water. Fortunately, Greg has provided me with a comprehensive bibliography of control systems theory and I can see that I've got a lot of catching up to do. My interests however are not, I think, incompatable with those of this group and I look forward to contributing to the upcoming discussions.

Some of the recent discussions pointed to a potential mapping between the Beer's and Perceptual control implementations. Thanks to Beer's book and Pat And Greg's NCSK program, I hopefully have some idea of where this group is heading and the direction looks promising.

Let me attempt to open up a fruitful debate by saying that whereas the central thesis of Beer's book was "Intelligence as ADAPTIVE behavior", my own interests over the past several years has been in looking at "Intelligence as EMERGENT behavior". By this I mean that intelligence can be looked at as an emergent behavior at the global level which arises from the collective adaptivity of a group of computational elements at a local level to external events.

While P. computatrix (Beer's cockroach) exhibited several fascinating (global) behaviors (wandering, edge following, food location) with a relatively simple neural control (local) structure, I would argue that these behaviors were not emergent as they were behaviors that Dr. Beer had explicitly "hard-wired" into the neural control system of his artificial insect. Thus, by my definition did not constitute acts of "intelligence". However, Beer did point in the right direction - A constructionist approach with attention to biological details is important to the development of any artificially intelligent agent. Local changes, his control system showed, can produce global behaviors. As an aside, let me say that prior to reading Beer's book, I had been following my interests from an Artificial Life perspective. There researchers took essentially arbitrarily constructed three layer neural networks, encoded their structures as bit strings, placed 65,536 on the nodes of a Connection Machine and mutated the lot for 500 generations to see what sorts of interesting behaviors emerged. All they got for their efforts were wee beasties that could walk, find food and bring it back to the nest -- and this was based on how they biased their fitness functions. It sure sounds like a brute force attack to me now. Greg got much more interesting behaviors out of his bug on a PC.

Ok. So Beer's insect produced biologically interesting behaviors without being intelligent. What's needed to generate intelligent (emergent) behavior? (e.g. If we put two of Beer's bugs together would they exhibit competative/cooperative behaviors? Would a language emerge? Why? Why Not? The answer to these question is where my interests lie. While many have assumed that the neuron is the fundamental unit of information processing in biological organisms there is growing evidence [1],[2] that the fundamental unit may be much lower. Is this the hierarchy in Hierarchical Control Theory? Also, while computer simulations are very fertile starting ground for any study of intelligent behavior, my own bias is towards a physical implementation of any theory. Studying AI without robotics is like studying General Relativity and assuming that the universe is void of all matter (You get pretty, closed form solutions to GR that way, but its a long way from reality). Anyone working on a hardware version of Beer's Bug or some PCT/HCT version thereof?

Alan E. Scrivner
c3141aes@mercury.nwac.sea06.navy.mil
----
[1] Rasmussen, Steen, et.al., Computational Connectionism within neurons: A model of Cytoskeletal automata subserving neural networks. Physica D, 42 (1990), pgs. 428-449. Elsever Science Publishers B.V. North-Holland.

[2] Langton, Chris G. Computation at the edge of chaos: Phase transitions and emergent computation. Physica D, 42 (1990), pgs. 12-37. Elsever Science Publishing B.V. North-Holland.

Date: Mon Feb 10, 1992 10:32 am PST
Subject: Re social control and Closed Loop 2.1

[From: Bruce Nevin (Mon 920110 12:48:50)]

(Robert K. Clark (Mon, 10 Feb 1992 15:38:00 GMT)) --

Not included in the recent issue of Closed Loop was the observation, by several people, that there is no demonstrated means by which one living control system can reach over and set the value of a reference signal in another. Also not included is discussion of the differences between control and social influence. This is in archives on the list server, but being told that is not very helpful without a date to narrow down a search. This last issue of Closed Loop covers much more than the past three months: my mention of Ruth Benedict about half-way through dates from last June, and much of the first half was unfamiliar to me, perhaps because I don't remember, but I think because I had not read archives from before about last April or May. I think that clarification of social influence vs. control was just before the August CSG meeting.

I think Greg hinted at a second issue on social "control."

Various "groupware" systems have been around in experimental settings for a number of years, which support group discussions by representing structural relations among foci of discussion, points of agreement and disagreement. We don't have that, nor can the listserver provide it (some sort of indexing function would be a start). Closed Loop serves some of that sort of function for us just by playing back by us some of what we have said to one another, which we otherwise might not remember and would be unlikely to reread. I am finding it very valuable.

Bruce
bn@bbn.com

Date: Mon Feb 10, 1992 11:18 am PST
Subject: Pic;robotics

[From Bill Powers (920210.0900)]

RE: Ptr4 crib sheet

>If you aren't at a menu or in the editor, you can press the "?" key ...
>Also, the KEYS.TXT file lists all key functions.

The help screen helps some, but it's oriented largely toward describing what keys can be pressed. The KEYS.TXT file lists the functions by the keys, also. If you know what key you want to press and want to know what it will accomplish, this is fine. But if you know what you want to accomplish (the more usual case), and want to know what key to press, the organization is backward -- you have to search through all the key entries to see which one does the thing you want done. If you want to center the screen on the cursor, you have to read all the entries clear down to the third-to-last entry in KEYS to find that you press X.

In PCT terms, the user wants to accomplish something (the known part) and wishes to find the action that will accomplish it (the unknown part). So help material should be organized by the known part in order to look up the unknown part. PTR4 is far from the only program I've ever see that organizes help material backward. This is why I proposed the crib sheet. I'm working on one for myself, and will send it to you for consideration when it's reasonably complete.

>And there is a PostScript output file with keyboard templates.

Let them eat cake, eh?

I agree with you about Beer. Right now I'm not in a position to offer specific suggestions about delineating what he has done "right", but you guys go ahead and I'll catch up.

Bob Clark (920209.1358) --

>I had thought that this subject ("Behaviorism") had long since
>disappeared of its own internal contradictions.

Many of us have been told by various journal referees that behaviorism is dead. They seem to have a very narrow view of what constitutes behaviorism, because they quite freely offer "corrections" implying that behavior is caused by what happens to an organism, and that behavior is simply the terminal event in a regular output process.

>The organizing principle of the Hierarchical structure provides a
>working framework for the analysis of complex behavioral systems.

It's hard to get this message across when the listener thinks that current beliefs explain adequately what is going on. This is why I like demonstrations -- they show that something is going on that current beliefs don't even acknowledge.

Kent McClelland (920209.1724) --

Re: emotions arising from wants.

>the real issues is where the "wants" come from.

I agree, but a secondary issue is whether the felt emotion causes the want, or is a consequence of the wanting. This is a rather old issue, which has been presented this way: do you run because you're afraid, or are you afraid because you run? Does the sight of the lion cause you to feel fear, which then causes you to run? Or does the sight of the lion cause you to run, which makes you feel afraid?

Both of these alternatives propose a simple causal picture in which the sight of the lion starts everything, leaving out your reference signal for seeing a lion. CT denies both. The sight of the lion means nothing unless you have an intention with respect to which you can evaluate the sight. You may tune in an animal documentary or visit the zoo because you want to see a lion. If you see a loose lion not in a zoo, you may decide you want to not see a lion, or rather, to see a tiny lion much farther away. As you are seeing a big lion up close, this results in a large error signal, which manifests itself in actions that reduce the amount of lionness seen, and also initiate a change in body state that is felt (in context) as fear. You get the hell out of there.

Where did the want come from? Not from the emotion, because the same emotion is involved in many situations, yet in other contexts it doesn't lead to running away from a lion. Not from the action, either, because running away from a lion is very much like running toward a rock star to get an autograph: it's just running. It comes from a goal for a perception, and from the fact that the perception doesn't match the goal. BOTH the action AND the feeling arise from this discrepancy, as part of controlling perceptions.

>Another student averred that emotions have a more central role in
>guiding behavior than you seem to allow. They are, she suggested, like
>a "12th level" of the hierarchy that provides reference signals for the
>levels below.

Is this student treating emotions as an independent force that acts on the nervous system? This is a very common theory, but I think it's an illusion. We often feel emotions without understanding where they came from, so it makes sense to think of them as controlling what we do. But counselling, therapy, or deeper reflection would, I think, uncover the want that lies behind any emotion. All one has to do is ask, when in the apparent grip of an emotion, "What is it that I'd like to do if I could?" If you're angry, what you'd like to do is to take some very aggressive and violent action to make someone be different or do differently, isn't it? There's some kind of big error, isn't there? If you didn't have that goal, if you didn't feel it was very important to oppose that person or stop whatever that person is doing, there wouldn't be any emotion.

The "emotion-as-cause" theory has long been used by men against women, women being reputed to have stronger emotions than men and being less able to think rationally when emotion takes over. This has been an effective way of minimizing the reasons for women's wanting to change the way they are treated and what they are allowed to do. If we look on the emotionality of women as a measure of the errors they are experiencing and the importance to them of the goals that are being denied, the picture changes entirely. They are not victims of emotion: they are experiencing enormous discrepancies between what they want and what they are able to achieve. Any human being experiencing an equal discrepancy would gird for action or prepare for escape, in proportion to the amount of suffering.

>Maybe we might consider emotions as part of the intrinsic error
>regulation system. Emotions might be perceptions that bring news of
>intrinsic error or its resolution by successful reorganization.

Intrinsic error implies a deviation of the state of the body from a reference state specified in the genes. So if the somatic component of emotion-perceptions arises from states of the body, your suggestion makes perfect sense: we should feel the variables that are, by their changes, resulting in intrinsic error signals. The emotions don't arise from reorganization, though. They arise from the same conditions that initiate reorganization. I think it's going too far to say that the bodily states we consciously feel are identical to the states monitored by the reorganizing system. There may be some overlap, but what we experience of the body gets organized into perceptions as the hierarchy grows, while the reorganizing system is monitoring intrinsic variables from the start.

On the other hand, as you say,

>Newborn babies certainly seem to come equipped with emotions, especially
>the negative ones like rage, fear, and general discomfort.

I think that these are adult interpretations: if I were behaving like that, I would be feeling rage and fear. However, the signs are certainly there that the baby can detect "wrong" physical states, and reacts to them by attempts at violent actions (big error). This probably entails sensing somatic states similar to those in an adult. But the baby probably isn't capable yet of having goals like "attack" or "flee," because those imply perception of relationships. So anger and fear in a baby wouldn't feel the same as they do to an adult -- they wouldn't be attached to such complex goals.

Re: take the money and run.

If we take the money, they'll just want it back later after we've spent it. Better to straighten it out now if possible.

Gary Cziko (920209.2040) --

Your Stage 4 coincides with my Big Picture. As to getting reorganization into models, this is something that takes a while to develop. First, I think, people have to get used to modeling control behavior and become comfortable with this as a procedure. I wish that most of the people on this net were at least trying out simple models and seeing for themselves how incredibly well they predict, instead of just reading about it. Those who have done this have, in fact, started to explore reorganization as well -- Rick Marken and Tom Bourbon, for example.

There's some conflict here. When you try to model something fairly difficult, like the behavior of gatherings or the details of hand-eye coordination, most of the effort goes into just finding a model that works properly. The reorganization is going on in the modeler, not in the model. When the model reaches some sufficient level of competence, you might then think about incorporating reorganization. But it takes a long time to get a complicated model to that level. I plan to look at reorganization in the arm model, but I'm still feeling bogged down with getting version 2 of the program into some sort of shape for release -- I haven't wanted to work on it for weeks -- and it isn't likely to include reorganization until that hurdle is passed.

Potentially, modeling reorganization is the most interesting aspect of a control-system model. It's necessary to decide what sort of parts-kit to start with, what variables to define as "intrinsic," and how those variables should depend on the behavior of the system. Each component has to have parameters subject to random variation by the reorganizing process, and those have to be selected. Perhaps interconnections, too, have to be made subject to reorganization, so one has to decide just how connections can be altered: what the possibilities and limits are.

When you've done all this, of course, you turn on the system and watch it develop into a hierarchy of control systems.

The most horrendous part of this project is simulating the environment. I think this part is so daunting that we will be forced to use the real environment, with real sensors and effectors. There's just no way to get sufficient richness and realism into a computer simulation of the environment. That will require more money than I will ever be able to command. Maybe Martin Taylor is in the best position here. Or maybe there are silent listeners out there who are headed in this direction.

In the meantime, we can learn something by trying to model elementary examples of reorganization. When we can.

Oded Maler (920210.0203) --

>Some people ... are trying to *build* working artificial systems ...
>For them, if they succeed in making something that works,
>the question what does it "really" control is not relevant. You can
>poke electrodes into their (creatures') sensors and tell them "your
>creatures really controls for this and that perceptual variable", maybe
>you can do it even directly by looking into their circuits and
>programs. So what?

This is a correct assessment of how they often react. What we have to make sure to communicate is that this isn't just a matter of definition. We have to persuade them to make their environments more realistic by introducing unpredictable disturbances that affect the outputs of their systems. They avoid doing this now, because the conventional approach assumes that if the output is disturbed the behavior pattern will be disturbed. That's why we see these ponderous machines designed so that output does exactly what is commanded, with normal disturbances being too weak to affect the output or else simply being prevented by the experimenter. Disturbances are thought of now as something that has to be sensed at the source and compensated for by some smart program that alters output commands -- or else as something to be prevented.

Consider what Greg and Pat Williams said yesterday, playing Devil's Advocate: the bug's antenna is stimulated by contact with the wall, which CAUSES the bug to veer away from the wall. This Gestalt has to be broken before the real application of control theory can be understood. One has to consider the whole circle: the bug's motion CAUSED the anntenna to be stimulated, which CAUSED the change of direction, which CAUSED the antenna signal to become zero again...

If we realize that this is a control system, we can then see that the antenna signal is the central object of control, not the bug's position. We could test this easily, by altering the position of the obstacle continuously, or bringing in any extraneous event that could stimulate the antenna. The position of the bug would change as required to keep the antenna signal at zero. This might then lead someone to think of changing the reference signal, which would result in the bug's *seeking* contact rather than avoiding it -- all without any change in the rest of the circuit. Through all these changes there would be only one constant relationship maintained over the long run by the system: antenna signal equal to reference signal.

This is why it's necessary to have some degree of realistic physics in the model. If you don't have the physics in it, you can't try out new kinds of disturbances to see how the model reacts to them. You can't see that the outputs vary all over the place while maintaining the INPUT at the reference level. I would like to test Beer's model by applying a sideward force to the body of the bug and seeing if its leg forces would change to oppose the effect of the disturbance (as I'm sure would happen in the real cockroach). But from what I know now, this wouldn't seem possible. When I know more about the model, maybe I will be proven wrong.

When you think in terms of controlling perceptual signals, the whole task of building a control system takes on a different appearance. For example, think of the six-legged machines going over rough terrain (very clumsily). If you watch a large animal walking on rough terrain, you see that the body follows a smooth path while the legs adapt to the varying footing. So clearly, height of the body above the (average) ground level is being controlled. Instead of trying to achieve this effect by sensing the terrain and calculating compensations for all the legs -- a very complicated method that doesn't work very well -- why not just start with sensing the variable that is to be controlled, height of the body above the ground? This information might be derivable from joint angles, or simply by a few sonic altimeters (if all you want is a model that works). The problem then is greatly simplified: if the sensed height deviates from the reference height, the error signal operates the various joints in the direction that keeps the error small. The problem boils down to sensing height above the ground, averaging in the way that results in a path of the model's body that matches the path of the real body.

>If you keep in mind the distinction between the two questions, I think
>you'll get more positive feed-back ...

Oded! Shame on you! We want our persuasions to move people toward our goals (negative feedback) instead of away from them (positive feedback).

Jean-Arcady Meyer (020210.0542) --

This looks like a conference at which CSG people ought to be represented. Far too expensive for me, but perhaps others could get some institutional support. It would be lovely to show an arm model, the Gatherings model, Rick's spread-sheet multilevel system, Tom's multi-person model, and maybe the Demos, at least to spread around a little. At least half of the subjects mentioned are pertinent to CSG interests.

Date: Mon Feb 10, 1992 11:29 am PST
Subject: Re: BEHAVIORISM

[FRom Chris Malcolm]

Bob Clark writes:

> BEHAVIORISM, DETERMINISM & FREE WILL

> There are Systems satisfying the criteria
> defining Control Systems that are composed of hardware. In such a
> non-living system it may be possible to trace all the signals/events
> through the system, so that piece-wise it demonstrates
> "straight-through" cause and effect (S--->R) operation.

Not necessarily. What about Watt's steam engine governor? A negative feeback control system built entirely of brass and steel, with a user-alterable set-point.

> It appears that the further developments in physics, Einsteinian
> Relativity, Quantum Mechanics, the Heisenberg "Uncertainty Principle"
> etc were not included in such deterministic explanations. In
> addition, the intrinsic impossibility of gathering enough data to
> predict behavioral responses seems to have been ignored.

Chaotic systems too. But note that the crucial point about determinism ("everything is pre-ordained") has nothing to do with the practicality

of forecasting. If forecasting is impossible that just means nobody can
predict what going to happen; whether or not the outcome is in fact
pre-ordained is another question enitrely. Computer simualtions of
chaotic systems provide a very nice illustration -- they demonstrate
that very simple physical systems can be constructed which are both
ineluctably determined and impossible in principle to predict.

> I have also been intrigued that the determinists seem to omit
> themselves from their own theory! They experiment, learn (are
> "conditioned"), and so on -- they certainly act as though they have
> some kind of Free Will!

Quite so, but I can certainly imagine a world in which I ought to behave
as though all of us had Free Will even supposing that we actually don't.
(I don't know whether our world is like this or not, but I hope to find
out!)

Chris Malcolm

Date: Mon Feb 10, 1992 11:50 am PST
Subject: Re: A.S. intro

From Greg Williams (920210)

>Alan E. Scrivner, 920210

I'm glad to see you're on CSGnet, Alan -- and bearing food for thought!

>... intelligence can be looked at as an emergent behavior at the global level
>which arises from the collective adaptivity of a group of computational
>elements at a local level to external events.

>   While P. computatrix (Beer's cockroach) exhibited several
>fascinating (global) behaviors (wandering, edge following, food location)
>with a relatively simple neural control (local) structure, I would argue that
>these behaviors were not emergent as they were behaviors that Dr. Beer had
>explicitly "hard-wired" into the neural control system of his artificial
>insect. Thus, by my definition did not constitute acts of "intelligence".

I would put it slightly differently: the "intelligence" reflected in the bug's
adaptive behavior is almost completely dependent on the (truly emergent, in
your terms) intelligence of Beer-as-modeler. I agree that reducing such
dependencies is an important aim for behavioral simulation modelers. As you
know, it's difficult!

>   Ok. So Beer's insect produced biologically interesting behaviors
>without being intelligent. What's needed to generate intelligent (emergent)
>behavior?

Maybe the control-theory modelers DO need NSCK version 4 -- the one with the
modifiable synapses. Several threads on the net seem to be converging at this
time on issues of reorganization. From the neuroethological point of view,
the concept of reorganization raises several (I think) interesting questions,
the first being "How can one tell when a non-human animal is reorganizing?" To
make it concrete: "what sorts of experiments would one do to decide whether or
not cockroaches reorganize?"

>Anyone working on a hardware version of Beer's Bug or some PCT/HCT version
>thereof?

Dr. Beer has a robot (shown in the Sept.-Oct. 1991 AMERICAN SCIENTIST) which incorporates gait control and (not in the original bug model) escape maneuvers. Gary Cziko might have more to tell on this. (And a question to Gary: Did Dr. Beer express any attitude at all regarding PCT when you saw him?)

Again, good to converse with you, Alan!

Greg

Date: Mon Feb 10, 1992 12:51 pm PST
Subject: spreadsheet model

Is the revised version of the spreadsheet model available in loadable ASCII? By revised, I mean with the changes in how weighting is done. I believe I now have means for loading 1-2-3 ASCII format into LUCID. Thanks.

Date: Mon Feb 10, 1992 1:13 pm PST
Subject: Beer and PCT; Reorganization

[from Gary Cziko 920210.1430]

Greg Williams (920210) writes:

>(And a question to
>Gary: Did Dr. Beer express any attitude at all regarding PCT when you saw
>him?)

The closest we got to PCT was my attempt to explain how PCT conceives of behavior as resulting from a hierarchy of control systems. I showed about 15 seconds of part of Demo1, 20 seconds of part of Demo2, and 13 seconds of the arm demo.

This led to his comment that I already posted about how he doesn't see why it has to be all one way (e.g., feedback control) or the other (i.e., central pattern generators) but rather an interaction of the two (e.g., hi speed gaits internally generated, low-speed gaits need sensory feedback).

Concerning the current discussion on reorganization, let me remind you guys of Beer's use of "off-line" genetic algorithms to come up with nervous systems that accomplish certain behaviors (or control perceptions). If this were made "on-line" within the bug, wouldn't this be a nice example of simulated reorganization?--Gary

Date: Mon Feb 10, 1992 2:45 pm PST
From: Dag Forssell / MCI ID: 474-2580

Subject: Teacher & Literature List

This is a follow up to my post of January 21.

I am preparing two lists, one of teachers and one of literature.

Teacher List:

Tom Bourbon, Professor *
Department of Psychology
Stephen F. Austin State University
Nacogdoches, Texas 75962

Gary A. Cziko, Associate Professor
Educational Psychology
University of Illinois
1310 Sixth Street
Champaign, Illinois, 61820-6990

Edward E. Ford, MSW, Lecturer *
Arizona State University
10209 North 56th Street
Scottsdale, Arizona 85253

Hugh Gibbons, Professor *
Franklin Pierce Law Center
Concord, New Hampshire 03301

Wayne Hershberger, Professor *
Department of Psychology
Northern Illinois University
DeKalb, Illinois 60115

Kent McClelland, Associate Professor *
Department of Sociology
Grinnell College
Grinnell, Iowa 50112-0810

David M. McCord Assistant Professor
Department of Psychology
Western Carolina University
Cullowhee, North Carolina 28723

Clark McPhail, Professor *
Department of Sociology
University of Illinois
Urbana, Illinois 61801

Hugh G. Petrie, Dean *
Graduate School of Education
State University of New York
Buffalo, New York 14260

Richard J. Robertson, Professor *
Department of Psychology
Northeastern Illinois University
Chicago, Illinois 60625

Philip J. Runkel, Professor Emeritus of *
Education and Psychology
University of Oregon
5070 Fox Hollow Road
Eugene, Oregon 97405-4008

Charles Tucker, Professor *
Department of Sociology
University of South Carolina
Columbia, South Carolina 29208

* Denotes author, editor, contributor. See literature list.

Literature List:

Powers, William T., BEHAVIOR: THE CONTROL OF PERCEPTION. 296 pages. Aldine DeGruyter Hawthorne, NY. 1973.

Robertson, Richard J., and Powers, William T. editors, INTRODUCTION TO MODERN PSYCHOLOGY; The Control Theory view. 238 pages. The Control Systems Group, Gravel Switch, KY. (1985, 1989) 1990.

Powers, William T., LIVING CONTROL SYSTEMS: Selected Papers of William T. Powers. 300 pages. The Control Systems Group, Gravel Switch, KY. 1989.

Marken, Richard S. editor, PURPOSEFUL BEHAVIOR: The Control Theory approach. American Behavioral Scientist, special issue. 11 articles, 16 contributors, 121 pages. Vol. 34 / Number 1 September/October 1990. Sage Publications, Thousand Oaks, CA.

Hershberger, Wayne, editor, VOLITIONAL ACTION, CONATION AND CONTROL. Advances in Psychology 62. 25 chapters, 33 contributors, 572 pages. North Holland, NY. 1989

Ford, Edward E., FREEDOM FROM STRESS. 184 pages. Brandt Publishing, Scottsdale AZ. 1989

Gibbons, Hugh, THE DEATH OF JEFFREY STAPLETON. Exploring the way Lawyers think. 197 pages. Franklin Pierce Law Center, Concord NH. 1990.

McClelland, Kent, PERCEPTUAL CONTROL AND SOCIOLOGICAL THEORY. Paper, 1991.

McPhail, Clark, THE MYTH OF THE MADDING CROWD. 265 pages. Aldine de Gruyter, Hawthorne, NY. 1990

Petrie, Hugh G., DILEMMA OF ENQUIRY AND LEARNING. ??? pages. Univ of Chicago press, 1981

Richardson, George P., FEEDBACK THOUGHT IN SOCIAL SCIENCE AND
SYSTEMS THEORY. 374 pages. Univ of Pennsylvania Press. 1991.

Runkel, Philip J., CASTING NETS AND TESTING SPECIMENS; Two grand
methods of psychology. 216 pages. Praeger, N.Y. NY. 1990.

As noted in the first posting, my purpose is to have handouts which convey the credibility of PCT which the committments and research by professionals to date give it. This is not intended to be exhaustive, but to provide a good start on a reading list.

If there are any mistakes, opportunities for clarification, or whatever - I will appreciate a note.

Dag Forssell
23903 Via Flamenco
Valencia, Ca 91355-2808
Phone (805) 254-1195   Fax (805) 254-7956
Internet: 0004742580@MCIMAIL.COM

Date: Mon Feb 10, 1992 3:31 pm PST
Subject: Re: PT; simple models; PCT's value in b.s.

From Pat & Greg Williams (920210-2)

>Bill Powers (920210.0900)

>This is why I proposed the crib sheet. I'm working on one for myself, and will
>send it to you for consideration when it's reasonably complete.

Thank you. We will try to incorporate it into PT if anyone besides you ever says he/she has your kind of problem with the help screen as it is now. Our beta testers have expressed satisfaction with the help screen as-is, but that doesn't mean it couldn't be improved. At least we put all of the info in one place, rather than hiding it in hierarchically structured menus (or so some users enthuse).

>I wish that most of the people on this net were at least trying out simple
>models and seeing for themselves how incredibly well they predict, instead of
>just reading about it.

We are all for trying out models, too. We think it is important to spell out explicitly what simple PCT models are capable of, to date, to counter those who say, i.e., that there is little or no empirical basis for Bill's hierarchical mechanism's details (assuming that the simple models which predict incredibly well incorporate some of that mechanism's details). And it would also be important to enumerate the limitations of the simple models' predictions, so nonPCTers don't get the mistaken impression that somebody is claiming the moon while grabbing at a reflection in the pond.

Bill, your comments on Oded's post provide an excellent introductory tutorial for behavioral simulation modelers which show the potential advantages of adopting a PCT viewpoint. We urge you to make these or similar comments available to Dr. Beer.

Kicking back on a beach in Hawaii is an enticing reference signal! Maybe we could re-site this year's CSG meeting? Only kidding, unfortunately.

Pat & Greg

Date: Mon Feb 10, 1992 4:20 pm PST
Subject: emotion;Big Pic;robotics

[Martin Taylor 920210 1845]
(Bill Powers 920210 0900]
>
>On the other hand, as you say,
>
>>Newborn babies certainly seem to come equipped with emotions, especially
>>the negative ones like rage, fear, and general discomfort.
>
>I think that these are adult interpretations: if I were behaving like
>that, I would be feeling rage and fear. However, the signs are certainly
>there that the baby can detect "wrong" physical states, and reacts to
>them by attempts at violent actions (big error). This probably entails
>sensing somatic states similar to those in an adult. But the baby
>probably isn't capable yet of having goals like "attack" or "flee,"
>because those imply perception of relationships. So anger and fear in a
>baby wouldn't feel the same as they do to an adult -- they wouldn't be
>attached to such complex goals.
>

I believe that if you inject someone with adrenalin, they have "funny feelings" but can't identify whether the feeling is elation, fear, anger, or what, until they see something that would account for the feeling (a nearby lion, perhaps?). This suggests to me that the somatic sensations are not the same as the identified emotions, and that to identify an emotion is a refined result of learning to associate internally and externally derived sensations at some high level, probably a level that involves some high-level control systems.

Bill, in browsing around for other things today, I came across a paper I had published with two colleagues in 1973. I thought it might amuse you so I am sending you a copy. In it, we discuss the difference in the sensations of being touched and of feeling (actively) an object. Only with active, internally controlled touch is there a perception of objectness. In the paper, we ascribed this difference to the action of a three-level hierarchic control system, in which the feedback kinaesthetic perceptions were melded with the sensations provided by the contact between the object and the skin. When the control hierarchy is inactive, we said, there is no perception of objectness. Behaviour as control of perception, with a vengeance!

The thesis of the paper, now I come to think of it, is just the same as the one I proposed above in respect of emotion--the combination of externally and internally derived sensation to form a coherent perception (which presumably forms the input to some ECS, according to the current models).

I had completely forgotten about this paper and the thinking that led up to it, but one of my co-authors is currently involved in robotics work, providing the robots with haptic perception. Maybe she remembers it better than I did. Anyway, this is a second reason why your ideas fell on such

fertile ground in my head. They were reanimating old images that had
never been allowed to flourish! Thank you.

Martin

Date: Tue Feb 11, 1992 4:38 am PST
Subject: Beer, PCT & CSGnet

[from Gary Cziko 920210.2040]

Pat & Greg Williams (920210-2):

>Bill, your comments on Oded's post provide an excellent introductory tutorial
>for behavioral simulation modelers which show the potential advantages of
>adopting a PCT viewpoint. We urge you to make these or similar comments
>available to Dr. Beer.

Another approach, Greg, would be for you to stitch together the relevant Beer Bug posts into a thread to send to Beer. You will probably wind up doing this anyway for _Closed Loop_, so why not a "special edition" for Beer that would make us look as friendly as we really are (just be sure to be careful with anything from Rick Marken, especially the "relatively nuts" stuff and the "this-is-kindergarten-stuff" comments).

Since he knows that you and Pat are intimately familiar with his first bug, he might appreciate your introduction to PCT more than someone else's from CSGnet.--Gary

P.S. I will try to get the log files to you as soon as they are ready to keep you up to date. Did you get log9202a?

Gary A. Cziko Telephone: (217) 333-4382

Date: Tue Feb 11, 1992 4:38 am PST
Subject: Bonnie Blair and CSGnet

[from Gary Cziko 920210.2145]

What does the winner of the first U.S. gold medal at the 1992 Winter Olympics and CSGnet have in common?

They are both based in Champaign, Illinois!

Champaign also has more CSGnet participates than any other locale and had the most contributors to the last issue of _Closed Loop_ (me, Joel Judd, and Mark Olson). I guess one could say that central Illinois is on the FAST TRACK in more ways than one!

Watch for Bonnie Blair Wednesday in the 1500 meters and Friday in the 1000.

Gotta run now and sharpen my skates.--Gary

Date: Tue Feb 11, 1992 9:37 am PST
Subject: Many subjects

[From Bill Powers (920211.0800)]

Joel Judd (920210.0856) --

>There are those who reorganize because there is intrinsic error arising
>from their attempts to reach their goals, and there is intrinsic error
>because others are forcing these to deal with goals not their own.

Agreed. The latter occurs not because the goals are not their own but because the required actions cause error directly or by conflicting with actions needed to maintain one's own goal-states. The "forcing" aspect may itself involve intrinsic errors other than existence of long-term or extreme error in the hierarchy. If it involves starvation or thirst, other kinds of intrinsic variables may be involved.

Another thought. Perhaps the existence of some nonzero chronic level of hierarchical error results in the lowest overall intrinsic error. In other words, it's better to learn than to know.

It's also possible that different kinds of intrinsic errors are weighted differently as to their contribution to overall rate of reorganization, so we would reorganize faster for some than for others. These weightings might change with age.

Bob Clark (920210.0920) --

>Whatever happened to the Concepts of a Hierarchical Array of Control
>Systems? And to the Test of the Controlled Variable?

They got into the social-control discussions, but not the parts in Closed Loop. You're quite right to mention the Test as a way of seeing whether the putative social control system actually exists. But it has to be evaluated judiciously. If you violate a social custom and everyone at the dinner table frowns at you, you might conclude that you've disturbed a social control system, which is exerting counterpressure. But if you carry the test to completion, looking for the means of control and the sensory channel, you will find only N individual control systems, each disturbed by your action. Blindfold all the people and they won't disapprove when you eat with your fingers. The social control system will disappear.

In your generalized 4-port black box, rather than defining input 2 as affected only by a disturbance, I recommend saying it is affected by the state of some external variable, so the output of port 4 depends on the state of the variable at port 2. Now we can distinguish the case in which there is an external link from port 4 to the variable at port 2 from the case in which there is a link from an independent variable to the variable at port 2. This also lets you define a "disturbance that doesn't disturb" -- the independent variable may or may not affect the state of the variable at port 2, because the output from port 4 may change enough to cancel the effect the independent variable would have when acting alone. So you can distinguish the existence of a physical causal link from the kind of effect obtained by varying the cause.

Bruce Nevin (920210.0845) --

The "rebus" concept is great -- how come you always know the right word?

Could this have anything to do with your profession?

I have a reply for John Underwood, Andover, Massachusetts:

```
         TAKE         2      VESTYOURMENTS
          I           C

                               Yours truly,

                                         O
                                     o
                                   o
                                 *   *
                                  * *
                           *       *       *
                            *     * *     *
                             *   *   *   *
                              * *     * *
                               *       *
```

I thought of that logo in 5th grade and have wanted to use it somewhere ever since. The Os should be a squiggle of smoke.

It's interesting to me that decoding these puzzles bounces around among nonverbal perceptions (words and letters as objects, spatial relationships, left-to-right sequence) and *spoken* language as well as written language. The conversion from "C" to "see" is based on the way "C" *sounds*, as is the conversion from "2" to "to". But the sense of relationship is derived simply from "TAKE" and "I" considered as meaningless objects arranged in space. The relationship is perceived, then named, then discarded while the name is retained: "under." Then the written and imagined words are assembled in spoken form to yield "I under take" and the spoken word "undertake" leaps out. The alternate possibility, "Take over I" may yield "takeover", but doesn't make sense when followed by "I". Decoding these things can take a while, during which while one tries out many possibilities that don't work, like "I below take."

This "rebusing" seems to work at levels other than categories. It clearly works at the relationship level, although not in sentences. It also works at the sequence level -- consider "Add the ingredients in the order flour, milk, eggs, seasoning." The sequence has no name, but it's perceived.

At the logic level there's a fairly direct translation: "If the door's unlocked, walk in, otherwise wait for me in the car" =

```
                           shrug, trudge back to car, get in
       try the doorknob, X
                           twist, push open, "HOO hoo"
```

where "X" marks the contingency perception, "unlocked?".

Notice that in decoding this,

```
        GOING
     c Cc c c ccc             (going overseas)
        cCcc
```

... it's necessary to perceive and name not just "C" and "c" but the nonverbal category "c's".

These kid games tell us a lot about language and perception. Now all we have to do is figure out what.

I suggested once that "zeroing" and reference questions may relate to the persistence of perceptions in attention after they have been evoked. In a sentence like

> John' father said, after Jim had told Mary about seeing the Parthenon on the trip to Greece he won in the contest a year ago next March, he was jealous.

The references can be figured out, but you have to remind yourself what the last "he" refers to: the meaning of "John's father" has faded out pretty much by then. You may have to re-read or replay the sentence. The time element is more evident in spoken language. If something like time-delay figures into zeroing, it would be pretty hard to find a linguistic rule that would explain the effect.

Of course your ideas about imagination apply too, and there are probably other considerations, depending on what's being zeroed or referenced. Sometimes the reference (as you said once) is only a placeholder, a variable whose value remains to be filled in, as in "what's being zeroed." It seems to me that this takes us to the program level where abstract symbols are manipulated only with regard to their position in a logical function and not their denotations.

I like what's coming out of this discussion, even though you probably see more implications in it than I do.

Martin Taylor --

Thanks for private comments. Pretty soon let's get that economics thread going. Control theory has a lot more to say about "demand" than economists have figured out. Think about it.

Alan Scrivner (920210.1053) --

Welcome to CSGnet! You will find much agreement here with your concept of intelligence as emergent behavior. You will also find some objections (interesting ones, I hope) to the idea that organisms adapt to external events. "Events" are only one of the kinds of things that organisms are concerned with, and they act just as often to adapt external things, changes, events, relationships, categories, and sequences to their own needs and intentions. We hope to repay you for your contributions with some concepts you may find suggestive. Argue all you like. We don't recognize existence of a category of questions called "dumb."

Best to all,   Bill P.

Date: Tue Feb 11, 1992 9:38 am PST
Subject: Beer and Behavior

[From Rick Marken (920210)]

Here is a first shot at a critique of Beer (and most of AI and robotics) from a PCT perspective. It became a bit longer and rambling than I had hoped, but I've no time to edit. I think this is a very worthwhile topic; hope there are some nuggests in the dust.

>Beer and Behavior

Beer is trying to make simulated organisms that behave adaptively. But what is behavior? In his book and in his American Scientist article Beer (like most everyone) uses words to describe behaviors -- "consummatory", "appetitive", "edge following" and "wandering", for example. But what do these words refer to? I propose the following definition of "behavior":

Behavior is variations in observable variables, the cause of which is attributed, by an observer, to an organism (or "agent").

Consider the behavior called "wandering". I believe that, for Beer, this word refers to variations in the 2 D coordinates (x,y position variables) of the bug over time. I believe that Beer would only count as behaviors those variations in x,y that could be attributed the bug itself. If the same variations in the bug's coordinates were a result of, say, wind pushing the bug around, they would not be called "wandering".

Now the question is, how does the observer decide whether or not variation of some variable should be attributed to an agent? Presumably by looking for the causes of these variations inside the agent itself (within the agent's "skin" -- which I take as a fair definition of the agent's envelope). There we find muscle tensions acting on joints that exert forces that produce the variations that we call "behavior". I will call the causes that we see as arising from inside the agent itself "actions". In Beers Bug, leg movements are the actions that produce "wandering" --ie variation in the 2D position of the bug. Actions may themselves be "behaviors" if they can be seen as having causes further inside the agent. But all we need to know are the "proximate" causes of any particular behavior. Thus, leg movements are the actions that cause wandering, neural impulses are the actions that cause leg movements. It is not necessary to determine the "ultimate" cause of behavior in order to see variation in the 2D position of the bug. Actions may themselves be "behaviors" if they can be seen as having causes further inside the agent. But all we need to know are the "proximate" causes of any particular behavior. Thus, leg movements are the actions that cause wandering, neural impulses are the actions that cause leg movements. It is not necessary to determine the "ultimate" cause of behavior in order to see variations in some variable as behavior --caused by the actions of an agent.

Now I can give a somewhat improved definition of behavior:

Behavior is variations in observable variables that are the result the actions of an agent.

I think this definition captures the meaning of the term "behavior" as it is used, in practice, in psychology and other behavioral sciences. Behavior is a "dependent variable" in psychological research. Typically, researchers measure dependent variables that they imagine to be a result of the actions of the agent under study. Example "dependent variables" include: time between stimulus and response, number of correctly checked boxes on a test, number of times a key is pecked/second, a spoken or written "response", volume of saliva in a tube (remember Pavlov), etc etc. None of these variables would be measurred as a "dependent variable" if the experimenter did not think that the agent under study had some causal influence on its value.

There are (at least) two problems with this approach to defining behavior.

1) The actions of the agent are rarely the sole cause of the variations

that we call "behavior" and (most important from my point of view)

2) The actions of the agent cause variations in many variables -- so ANY ONE of those variables can be legitimately called a "behavior".

Problem 2 is the reason I say that the conventional approach to defining behavior is ARBITRARY. Basically, the conventional approach says that "behavior" is any result of actions that I (the observer) find interesting. This is the point that I try to make with the "Mindreading" program. In that program, a subject uses a mouse to determine the 2D position of a number on the screen. Actually, the mouse influences the position of 5 numbers on the screen. The mouse is clearly the proximate cause of the position of the numbers. So the actions of the subject (mouse movements) have 5 simultaneous effects (ignoring stuff like noises and marks on the table produced by the mouse movements). The movement of ANY of these 5 numbers could be considered the subject's behavior. But the subject is asked to move only one number at a time. Only the movement of that number is "behavior" from the subject's perspective. If the subject is moving the "3" and you record the position of the "4" as a measure of behavior then your measure of behavior is WRONG -- at least from theÉ;ubject's perspective. However, it would certainly be counted as a legitimate measure of "behavior" by conventional psychology.

The "Mindreading" program, by the way, can determine which number the subject is intentionally moving -- that is, it can tell what the subject is actually "doing". It does this using control theory; it continuously performs the "test for the controlled variable" to all 5 numbers. Each number is influenced not only by the mouse but also by a slowly varying random disturbance -- a different disturbance to each number. The disturbance to the intentionally moved number will be systematically opposed by the subject -- yielding a fairly strong negative correlation between action and disturbance. Thus, the program detects the intentionally moved number by looking for the number with the largest negative correlation between mouse and disturbance. It works every time. I call it "Mindreading" because the program is set up so that it is impossible to tell which number is being moved intentionally just by looking at the numbers -- that is, just by looking at "behavior" as it is defined behaviorally.

Beer's approach to modeling suffers from the problems of the behavioristic definition of "behavior" given above -- particularly problem 2. He begins by defining (verbally) the "behaviors" that he wants the bug to perform. But he implicitly defines these "behaviors" conventionally -- in terms of variable results of actions. "Wandering" for example, is implicitly defined as "random changes in the 2D position of the bug". But it is actually implemented as "lateral foreleg-extensions" that occur "at random". The leg extensions are actions that have many results besides "random changes in the 2D position of the bug". It changes the position of the bug relative to an infinite number of points in 2 space, it changes the angle of the body with respect to an infinite number of different axes, etc etc.

The ambiguity of the behavioristic definition of behavior makes it difficult to design a system that behaves in a particular way. This is the problem with P. computatrix. The "random changes in the 2D position of the bug" is a behavior of the bug in the same sense that movement of any one of the numbers is a "behavior" of the subject in the "Mindreading" demo. Both are results of actions -- but not necessarily controlled results. In fact, 2D bug position is not a controlled result the way it is implemented. It is actually an action that is part of "edge following" -- which is (a least in part) a controlled result, because the bug is designed to keep its antenna sensor in a particular state. Parallel orientation to the edge may also be controlled (and influenced by leg actions) -- I can't tell from reading the American Scientist article.

The point is that, by looking at things in behavioristic terms, Beer makes the design process much more obscure and haphazard than it needs to be. This is especially true if he wants to design systems that will behave in disturbance prone enviroments (as I presume he does -- I think that is what he means by "adaptive"). His behavioristic view leads him to try to have the bug produce particular results of action. If he knows that these results involve "stimuli" (as in edge following -- the movement must be done with respect to a stimulus edge) then he adds a sensor. In fact, only when there are sensors for these "stimuli" is there control -- and "adaptation".

The control approach makes things much simpler and more elegant. Instead of defining behaviors as results of actions you define behaviors as CONTROLLED RESULTS OF ACTION. A controlled result is a perceptual result; so if you want the bug to control a particular variable you know that you need to design a way for the bug to perceive that variable; you also know that you need a way for the bug to affect the variable in all dimensions in which the variable is to be controlled.

If, for example, I want an "edge following" bug then I would first define "edge following" in terms of controlled perceptual variables. One might be "equal amount of excitation at two pressure sensors located on the side of the body". Then I design some action system (like the bug locomotion system) that can move the body so that the pressure sensor positions can change (output affects input). I can put in an explicit or defacto reference level for the sensor output (perception). The difference between reference and sensor output is the error that drives that locomotion system; this has to be set up so that increases in error move the legs in such a way that the sensor output is brought closer to the reference -- ie- negative feedback.

Beer has these kinds of loops in his model; it's just hard to see what the model is controlling and what is just uncontrolled action that is part of the influence on the controlled variables. I can understand why Beer wants to get the bug to behave in a particular way -- from the point of view of an observer. That's the whole point from the robot builder's perspective. But if you want an "adaptive" robot you have to do a little building from the robot's perspective -- you have to ask "what variables does the robot need to control for ITSELF in order to produce a particular observable behavior for me"?

Beer's bug does produce cute behavior. But it is not based on useful underlying principles. If the game here is to produce behavior that is fun to look at for an observer, then Beer's bug is fine. But I think for sheer amusement Bill's "Gatherings" program produces just as much or more interesting "emergent" behaviors -- and it's based on principles of control that are well understood. Moreover, Bill's "Little Man" demo shows (in beautiful detail) some of the complex, coordinated behavior that the bug shows -- and again, the basic principle of design is quite clear. Again, if the goal is to produce amusing "behavior" -- ie, results of action, then perhaps PCT could get some notice by building ping-pong playing robots; I bet it would not be too hard to give the "Little Man" another arm and have him juggle instead of just point.

I don't think these kinds of demos are really worth it because they are truly "superficial" -- as are all behavioral approaches to robotics.

It is the underlying principle that is important; if you want an organism that can produce consistent results in a disturbance prone environment then you must design it to CONTROL variables in that environment -- not to just cause variations in those variables.

One last point -- Beer's leg movement algorithm looks like fun. It looks to me like there is control involved -- because there are flexion sensors at each leg. The fact that the bug adjusts its gait to changes in extension in the front legs indicates that there is perceptual control going on -- and it is going on ALL THE TIME. Contrary to what Beer said to Gary Cziko, the bug (when a controlled variable is involved) is ALWAYS controlling. Changes in the rate of leg movement doesn't move the bug from open to closed loop; Beer just interpretes it that way because he is looking at the "behavior"; but the perceptions are ALWAYS there -- influenced by the outputs directly and indirectly. All speed might influence is how well these perceptions can be controlled. But we know that control systems have dynamic limits.

One more last,last point.

>On reconciliation.

My guess (and I hope I'm wrong) is that it will be as impossible to get Beer to see things from a PCT perspective as it will be to get an operant conditioner to see reinforcement as controlled input. Beer can obviously answer this on his own -- it's just my prediction. I hope I am wrong. But my experience is that coming close to PCT is no guarantee that one will cross the line. Operant conditioners are about as close as you can get -- my god, they are sitting there watching animals precisely control variables day in and day out. They watch them resisting disturbances to "set points" and they still don't (won't?) accept the idea that reinforcement is a controlled perceptual input (a concept that would require a new understanding of thir subject and a new approach to how they study it). I know (by hearsay) of only one behaviorist (T. Verhave) who actually got the point of PCT.

Beer does design circuits that are basically control systems. He currently thinks of them as S-R systems guiding "behavior". Whether or not he would see the PCT approach as something that could help him (which it would) I don't know. Maybe we could convince him that it might be worth it by building bugs that do what his do -- and then some -- using PCT. It would not really be hard. I guess what we need is for Beer to tell us what we could build that would convince him that PCT is worthwhile. If we pick it, I doubt that he would consider it interesting.

So how's that -- ask Beer if there is anything that is causing him difficulty in his approach. If he's got no error, he probably has no need to look into PCT anyway.

Sorry for the length of this.

Hasta Luego

Rick

Date: Tue Feb 11, 1992 10:01 am PST
Subject: Oded and Gary

from Ed Ford (920211.10:28)

Oded Maler - Closed Loop and CSG list is on the way via air mail.

Gary - Freedom From Stress not yet in Japanese. Two other books of mine are.

Ed Ford ATEDF@ASUVM.INRE.ASU.EDU
10209 N. 56th St., Scottsdale, Arizona 85253 Ph.602 991-4860

Date: Tue Feb 11, 1992 12:48 pm PST
Subject: rebates

[From: Bruce Nevin (Tue 920111 14:31:13)]

Joel, Gary: I notice credit given to Bill Williams in the annotated bibliography sent out to new CSGnet participants. Sure that isn't Greg Powers? ;-)

Bill Powers (920211.0800) --

Thanks, I thought you had used the term "rebus" and I was the one taking it up as a good term. But what's a little extra agreeability among friends?

In most cases, the basis for language control is in the spoken form, and the visual aspects of the written record are secondary. Some confusion arises due to the ambiguity of English orthography with respect to intonation. (All writing systems leave out information, usually "suprasegmental" information of intonation contours, often segmental information as well.) Thus, in your example for short-term memory:

> John's father said, after Jim had told Mary about seeing the Parthenon
> on the trip to Greece he won in the contest a year ago next March, he
> was jealous.

When spoken, the phrase within the bracketing commas has the reduced intonation of an interruption. The phrase "he was jealous" returns to the intonation contour of "John's father said he was jealous," at the pitch and volume appropriate for continuing that pattern as though it had never been interrupted. The paired commas represent the reduced intonation of the interruption, so you could say that English orthography covers that particular suprasegmental element here--with the caveat that not everyone uses English punctuation as consistently as they control English intonation patterns.

What would happen if the speaker was concerned that the interruption was too long, and the hearer or hearers no longer remembered what came bef∩rw the interruption, as often happens, as it is now--what would happen would be a reiteration of the first part (or the syntactic core of it, sans modifiers):

John's father said, after Jim had told Mary about seeing the Parthenon
on the trip to Greece he won in the contest a year ago next March--

he said he was jealous.

he said John was jealous.

John's father said he was jealous.

he said the boy was jealous.

The explicitness of the resumption depends on whether the speaker feels that the antecedents of the two occurrences of "he" are salient enough.

Still not clear whether/how a program ECS can "know" that two John/he perceptions refer to the same individual--that the rebus-link from the words to the category perception entails a thread of perceptual signals down to the same sensory input or imagined perceptions corresponding (as we take it) to one particular individual human being, for both words. That is the condition under which one of the words can have a reduced form (zero, pronoun, referential use of a classifier word like "boy"). Any help?

Bruce
bn@bbn.com

Date: Tue Feb 11, 1992 2:24 pm PST
Subject: CT Economics

[From Bill Powers (920211.1500)]

Here's a kickoff for a thread we might call CT Economics.

It's hard to get a feel for the antiquity of Adam Smith's writings in _The Wealth of Nations_*. But there's passage in my Penguin edition that helps:

> "In the province of New York, common labourers earn three shillings and sixpence currency, or two shillings sterling, a day..."

It isn't the wages or the strange dual monetary units to which I refer, but the footnote attached to "earn":

> "1. This was written in 1773, before the commencement of the late disturbances."

This is the book, as old as the United States, in which the "invisible hand," the Law of Supply and Demand, was invented. Economists have taken this law as a basic given of economics ever since. It's been taken as the magic formula that sets a free-enterprise system straight when it goes astray, no matter what mistakes are made. Leave the system alone and the invisible hand will make sure it flourishes. Meddle with it, regulate it, insert human intentions into it, and only disaster can follow. The Law of Supply and Demand is greater than any of us and operates in majestic disregard of mere human desires and purposes -- or so we are told by theoreticians in economics.

This, however, was not Adam Smith's conception of this law. For Smith, it was simply the natural consequence of the fact that people want things and are willing to devote some labor (or labour) to getting them. Here is what he said about demand.

> "The market price of every particular commodity is regulated by the proportion between the quantity which is actually brought to market, and the demand of those who are willing to pay the natural price of the commodity ... to bring it thither. Such people may be called the effectual demanders, and their demand the effectual demand; since it may be sufficient to effectuate the bringing of the commodity to market."

> " ... When the quantity of any commodity which is brought to market falls short of the effectual demand, all those who are willing to pay the whole value of the rent, wages, and profit, which must be paid in order to bring it thither, cannot be supplied with the quantity which they want. Rather than want it [lack it] altogether, some of them will be willing to give more. A competition will immediately begin among them, and the market price will rise more or less above the natural price, according as either the greatness of the deficiency, or the wealth and wanton luxury of the competitors, happen to animate more or less the eagerness of the competition. ... Hence the exorbitant price of the necessaries of life during the blockade of a town or a famine."

Similarly, he says, when there is an excess of supply over effectual demand, some of the goods must be sold to those who are willing only to pay less, so the market price will sink lower than the "natural" price.

"Effectual demand" is one side of a purely psychological explanation of the Law of Supply and Demand (the other side would describe the producer-seller's adjustments of the price, given the need to pay the costs of production and distribution and the desire to have a profit left over).

The payment of money for goods, Smith pointed out, misrepresents what is actually traded. "The real price of everything, what everything really costs to the man who wants to acquire it," he said, "is the toil and trouble of acquiring it."

So we have the basis for a model of the demand side of economic transactions. Human beings desire goods and services, some being desired because human beings can't live without them and the rest because human beings have learned to value them. People are willing to exert a certain amount of toil and trouble in order to bring the goods and services thither -- a great deal of toil and trouble, if they require the goods and services to say alive. If the shortfall is great, they will put out more effort to get what they want and need; if there is an excess, they will not make any effort to get more than they want or need.

On the supply side, there are people who also want goods and services, and are willing to go through toil and trouble to get them. Their labor is not on the production line, but is used to buy materials, rent facilities, borrow, organize, and otherwise manage production in order to create the goods and services that they and others want. They give receipts for goods and services to laborers who transform raw materials into production machinery and commodities, and they adjust prices (which are paid to them using the same receipts they gave out) for those commodities, so that the whole product can be sold and the costs of production can be repaid. The receipts returned to the producer-seller are used in part to pay the laborers who harvest the raw materials, build the machinery, and produce the goods and services; in part to repay the borrowed receipts with interest; and in part (the part they are most concerned about) to spend on goods and services for their own consumption.

So all human beings produce toil and endure trouble in order to bring the quantity of goods and services made available to them by that toil and trouble to the level that they need or want. This is the great engine that drives all human interactions, including those that come under the heading of economics. The invisible hand is the hand of a human being working to control what happens, not to an abstract economy, but to that human being. There are no other forces at work.

Theoretical economics has dropped the human being from these equations and has tried to explain the workings of an economic system with no human beings in it. This is why theoretical economics has so little relationship to what actually happens in this economy. In order to build a clear picture of economic interactions, we must understand that they result from the basic nature of living control systems, human beings.

Best

Bill P.

Date: Tue Feb 11, 1992 3:56 pm PST
Subject: Re: rebates

[Martin Taylor 920211 17:30]
(Bruce Nevin 920111 14:31)
>
>
>In most cases, the basis for language control is in the spoken form, and
>the visual aspects of the written record are secondary. Some confusion
>arises due to the ambiguity of English orthography with respect to
>intonation.

A bit far afield from CSG, perhaps, but I personally don't go along with this often expressed view. It is not at all clear, for example, that the original forms of written language had any connection whatever with the speech of the writers. Mostly the forms were related to visual impressions of concepts, such as grains of wheat, suns, rivers, or whatever. In cases of ambiguity, it was not uncommon to annotate one form with another, sometimes expressing visually a related concept, sometimes expressing a concept that had a spoken word that sounded like one of the possible spoken words for the ambiguous form.

Many Chinese characters and poly-character words are made in this way even now. It is only some 2500 years ago that any language was written in a way that could, with effort and skill, be transformed regularly into spoken words of the language. Even in languages with claims to be perfectly phonetic, there are often cases in which phonetic exactness is subordinated to visual clarity (either to disambiguate homophones or to maintain the relationship among families of words with common roots). Korean Hangul is a case in point--a written language deliberately designed to be a phonetic transcription of the spoken language, in which the written forms evolved so as to eliminate the visual similarities among same-sounding words, especially if they belonged to different families of words.

I prefer to see written language and spoken language as inter-related dialects in which literate people are bilingual. Each influences the other, but neither is truly primary. Our spoken language changes to conform to the writing, and vice-versa, but they diverge, too. As we discussed many moons ago, it is extremely difficult to make sense of literal transcripts of most spoken dialogues (and even monologues). There has to be a significant translation process before what was spoken can be placed on the page in such a way as to carry the same sense as that of the original speech.

Martin

Date: Tue Feb 11, 1992 4:16 pm PST
Subject: education

[from Joel Judd]

Bill, Gary, Bruce, Martin, Hugh (still there?)...

I have reencountered a paper by Muriel Saville-Troike that really read differently now that I'm trying to work out some research implications of PCT for learning. For several years she directed projects which would look at small numbers of individual children in bilingual environments, sometimes attaching a microphone to the child so that every single instance of language was captured for a certain amount of time, and compared with the videotaped context. I don't know if she has never been widely accepted among SLA researchers because she has dealt so much with children, or because she insists on maintaining the cultural context in language research, or both.

Anyway, one of the reports on three children from 1985 has some gems regarding the interpretation of schooling, culture, and the use of language. Two Japanese (M and F) and on Korean girl were followed during their Kindergarten year. Among other things, Saville-Troike reports how they became socialized in the elementary school--the "rules and directives of the teacher." The earliest English recorded dealt mostly with the many commands the children received ('put away your toys') and the Japanese boy was recorded saying these phrases to himself, and modifying them in directing the Japanese girl. Evidence of category development surfaced in another early production "That's mine." But even more telling was higher level development evidenced by the Korean girl's acquiescence when faced with the consequence "Then you not sit by me in the bus."

The children soon realized that they were dealing with another "language" when they realized that others did not understand them. They then would try to mimic L2 sounds (and other children would mimic them) realizing that the L2 was needed but lacking the skills to produce it correctly. Then all three fell silent for more than a month, and the Japanese girl never did progress in verbal English.

Because this school has pull-out language for the foreign students, they also had limited interaction with a native speaker of their L1 and to some extent its attendant culture. This provided some interesting contrasts. When the Korean girl transferred her "peer challenging" behavior from the Korean group to her K class, she alienated the students. If she cried in K class, the teacher comforted her; if she cried in Korean class, the teacher would become more "aloof" and tell her "the Korean equivalent of 'Don't be a baby.'" In another example the K teacher selected a Japanese boy to be the "leader" since his English was best. He began to translate and tutor other kids' English. Later when the Japanese teacher had to come in the class to discipline the Japanese group, she disciplined the boy in Saville-Troike's study because he was the oldest and therefore responsible for the group. According to Troike, this conflict engendered by the teachers alienated the two boys from each other for 3 weeks.

She mentions other cultural niceities which the children had to deal with in school. The kicker, though, comes when she asks the K teacher how the 3 children did during the year:

"It is interesting to note that although the teacher had devoted considerable time during the year to teaching knowledge about a variety of subjects and beginning skills for reading and writing, the criteria she used to judge relative readiness for first grade involved ONLY WHETHER RULES FOR APPROPRIATE BEHAVIOR HAD BEEN LEARNED AND WHETHER THE CHILDREN HAD DEVELOPED THE 'RIGHT ATTITUDES' ABOUT SCHOOL AND LEARNING. Whatever content knowledge was acquired during this socialization process was clearly a secondary consideration" (p.57).

To the extent that this teacher (or any similar teacher) can effectively judge such criteria, and Saville-Troike argues that she did (based on cases with other students), it would be interesting to know HOW she did it, if she was using some sort of "folk" knowledge of the Test. This also implies that the teacher is complicit in furthering the "right attitudes" fostered by those in charge of the school, etc.; that she agrees with such perceptions of schooling and has some form of them herself by which to judge student behavior. Some of her concluding comments provide more ideas for school research:

"The context of the school provides a conveniently CONSTRAINED environment in which to observe the workings of cultural input and intake...from explicit formulations mediated or regulated through language to implicit understandings communicated nonverbally...We must begin to escape the inherited tyranny of a focus on linguistic form, and even go beyond the limitations of communicative/pragmatic analysis, if we are ever to achieve a truly holistic perspective on the process by which children acquire a second language."

And from the same anthology, from Merril Swain, comes this conclusion about immersion programs which seem to not develop "native-speaker competence" in students:

"That is to say, the immersion students have developed, in the early grades, strategies for getting their meaning across which are adequate for the situation they find themselves in: they are understood by their teachers and peers. There appears to be little social or cognitive pressure to produce language that reflects more appropriately or precisely their intended meaning: there is no push to be more comprehensible than they already are. That is, there is no push for them to analyze further the grammar of the target language because their current output appears to succeed in conveying thier intended message."

Who wants to bother with the intricacies of another language when you've already got one, and the second is "good enough"?

Date: Tue Feb 11, 1992 4:41 pm PST
Subject: Re: Beer and Behavior

[Martin Taylor 920211 17:45]
(Rick Marken 920210)

Rick's new redescription of behaviourism led me once again to thinking about the topic, and what it must mean (quite apart from what any School has claimed for it). I take for granted that the data of any psychology must be the observation of what people (or animals) actually do. There isn't any other data except for the observation of what seems to go on in one's own head, and even if one is correct about that, one can't be sure that the same sorts of things are going on in someone else's head.

So. Behaviour is the measure of all things psychological. Wherein do PCT enthusiasts differ from the despised Behaviorists? PCT people believe that what people do is done because that will allow the people to set up desired perceptual conditions. Behaviourists (depending on the flavour) look the other way around the feedback loop and say that people do what they do because they want something (i.e. they say "the subject picked up the water glass and took a drink," perhaps adding "to assuage his thirst"). They say that taking a drink is more frequent under conditions of water deprivation, and stuff like that. What do PCT people say? The reference level that corresponds to the perception of thirst was lower than the perception, so the subject controlled the perception, in this case by drinking, which in this case involved picking up the glass, which in this case had water.

Whether you are a behaviourist, looking back up the different levels o abstraction of behaviour, or a PCTer looking in the signal direction up the different levels of abstraction of perception, you have much the same problem--to identify some property that remains stable against the variety of things that change in the circumstances of action. Both, in the end, come round to the idea that there are stable abstractions, and if the Powers model is correct, then both find stability in the same place: at the comparator of the control system, in the reference signal, or intention.

I think that most people would dismiss the behaviourists who say that behaviour is "only" the muscle twitches, just as we would dismiss those who would say that perception is "only" the responses of sensor systems. And just as there are an incredible variety of muscle actions that can accomplish the same intent, so there are multitudes of sensations that can lead to the perception that determines how well that intent is accomplished. The stability is internal, and inaccessible except by modelling (= theory) in either case.

The feedback loop involves a one-way information flow from intent to action to effect to perception (no, I'm not forgetting it is continuous). The observable things are in the range action to effect. I don't see any obvious reason why there should be a great benefit to following the flow forward or backward, if both routes arrive at the same place.

Martin

Date: Tue Feb 11, 1992 9:48 pm PST
Subject: Behavior

[From Rick Marken (920211)]

Yes, folks, I'm missing the Olympics (and that young lady from Illinois) just to make this reply to Martin Taylor. Boy, I hope I'm right.

Oh, I've just been informed that it's being taped -- so I can be wrong and not mind.

I want to begin by thanking Martin for making it perfectly clear that I failed to make myself perfectly clear.

Martin Taylor (920211 17:45) says:

>I take for granted that the data of any psychology must be the
>observation of what people (or animals) actually do.

But what is DOING?

>There isn't any other data except for the observation of what seems to go on
>in one's own head

What you see people doing is going on in your head. But that was not the point of my post. I was not arguing for introspection; I was arguing that the definition of behavior implied by ALL psychologies -- behaviorism, cognitivism, existentialism, Freudianism, Skinnerism, structuralism, functionalism, primal screamism, psychometrics, multidimensional scaling, signal detection theory, etc,etc -- is WRONG.

>So. Behaviour is the measure of all things psychological.

No problem. But WHAT IS BEHAVIOR??? If you say it is what people DO then what is it that people DO?

>
    Wherein do PCT
>enthusiasts differ from the despised Behaviorists?

We hate 'em all -- cognitivists even more than behaviorists. Harharr. (I'm wearing my Long John Silver Patch too).

>
 PCT people believe
>that what people do is done because that will allow the people to set up
>desired perceptual conditions.

I don't agree with this (I don't really understand it). What I (a PCT person) believe (and can demonstrate) is that what people do is control their perceptual experience. But this belief of mine is irrelevant to my argument. I wasn't trying to say what people see as the "causes" or the "explanations" of behavior; I was trying to say what people seem to have in mind when they say the word "behavior". WHAT IS BEHAVIOR?

>Behaviourists (depending on the flavour) look the other way around the
>feedback loop and say that people do what they do because they want
>something (i.e. they say "the subject picked up thewater glass and took a
>drink," perhaps adding "to assuage his thirst").

Again, irrelevant. It assumes that we know what "people do". If what people
do is so obvious, then one should never be able to look at a person's behavior
and ask "what the hell are they doing?" -- after all, one can SEE what the
person is doing, can't they? But people ask this about behavior all the time.
How could this be??? That is what the Mindreading demo is about. You
know that the subject is moving all 5 numbers -- no question about it -- yet
observers of the subject's behavior are often inclined to ask "what is that
person DOING?" ; the program answers the question automatically.

>Whether you are a behaviourist, looking back up the different levels o
>abstraction of behaviour, or a PCTer looking in the signal direction up
>the different levels of abstraction of perception, you have much the same
>problem--to identify some property that remains stable against the variety
>of things that change in the circumstances of action.

Ah, now we're getting somewhere. I think you know about CONTROL!
"Behavior" is a property (I would say a variable) that remains stable (I
would say in a possibly variable reference state) in the face of variable
disturbances. The problem is that, under ordinary circumstances, these
disturbances are invisible and the controlled variable is by no means
obvious. When a rat presses a bar to avoid shocks is it controlling the
probability of a shock or the interval between shocks (or some other
variable aspect of the shocks)? (see Powers' 1971 paper in Behavioral
Science). You can't tell by just looking. So if behavior is control (as I
argued
in my prior post and as you appear to concede in the above comment) then
you must test to determine what variable is being controlled; NO THEORY
NECESSARY HERE, JUST AN UNDERSTANDING OF THE NATURE OF CONTROL.

You seem to claim that behaviorists (and all other psychologists) determine
what an organism is controlling (DOING) as a matter of course (because they
study "behavior" and this implies that they know that behavior is control).
This is a very big surprise to me. If it's true it obviates a great deal of my
work in PCT. Heck, most of my work is oriented toward trying to show
psychologists that they take it for granted that they know what behavior is,
but don't. So if you know of studies where psychologists (of ANY stripe)
have done their work by testing for controlled variables, show me; please. I
would no longer be crying in the wilderness; indeed, the cry would be a false
alarm.

>
Both, in the end,
>come round to the idea that there are stable abstractions, and if the Powers
>model is correct, then both find stability in the same place: at the
>comparator of the control system, in the reference signal, or intention.

Again, I am amazed to find that all my work has been for nought.
Psychology already knows that behavior is controlled perceptual variables.
Shew!!

>I think that most people would dismiss the behaviourists who say that
>behaviour is "only" the muscle twitches,

Muscle twitches might be controlled too. It is not the behaviorists I have been dismissing. I am dismissing ALL psychologists who fail to recognize behavior as controlled consequences of action.

Look at it this way; do psychologists who use, say, "number of menu picks" as a response measure (you're a computer interface person too, right) test to see if this "behavior" is a controlled variable. Of course, not. They measure this behavior because it matters to THEM (the psychologists) for whatever reason. In general, psychologists pick "respense measures" because THEY THEMSELVES care about those measures; they don't ask themselves if this is something that the person is DOING (in the control theory sense; ie. CONTROLLING). It's a "behavior" and maybe it's related to other variables -- and off they go doing their "science". I am trying to show that this approach misses an important point; it neglects the fact that some behavior is intentional (CONTROLLED) and some is unintentional (an accidental side effect; uncontrolled; but a behavior nevertheless, in the sense that it is a result of the organisms actions). If you just measure the behaviors that are important to you, with no concept of whether or not they are important to the behaving agent, then you have no notion that the person might be controlling. I think it is a very good bet that, from the point of view of the behaving agent, much of the "behavior" that is measured in psychology is an ACCIDENTAL SIDE EFFECT OF ACTIONS. Even if this is not the case, it still means that a great deal of what passes for "bnehavioral" research may be completely irrelevent to our understanding of behavior.

>The stability is internal, and inaccessible except by modelling (= theory)
>in either case.

I am not trying to make a theoretical point here. I am trying to point out that psychologists have missed seeing a PHENOMENON -- the phenomenon of intentional (or purposeful) behavior. This is an objective phenomenon and it can be recognized without any theory (James did it over 100 years ago). But it is a hard phenomenon to study without being able to model it. Psychology could not model it (until the 1940s) so they said it didn't exist; behavior is just OUTPUT, like water from a spigot or light from a filament. I am trying to show that, by chickening out about purposeful behavior, psychology put itself in the position of having no basis (other than the arbitrary OPERATIONAL DEFINITION -- "behavior is the variable that I measure") for answering the question "WHAT ARE YOU DOING?".

>I don't see any obvious
>reason why there should be a great benefit to following the flow forward
>or backward, if both routes arrive at the same place.

Again, I don't think this is a theoretical question. My problems with what I (perhaps restrictively) refer to as a "behavioral" definition of behavior would exist without PCT (though PCT makes it all easier to understand). The Mindreading program (which I dare any psychologist to experience and remain a conventional psychologist) is not based on theory; it is based on an understanding of the PHENOMENON of control. There is no way to EXPLAIN the results of the Mindreading demo, however, without PCT.

Best regards

Rick

Date: Tue Feb 11, 1992 10:14 pm PST

Subject: Re: emotion; behaviorism. Yngve

On EMOTION: For anyone interested in an account of emotion as felt states of the body that accompany action, I recommend a little book by Andy Papanicolaou, who attended tha last meeting of CSG. It was written shortly before Andy took the plunge into PCT and I think you will see why that plunge was not very big.

Andrew C. Papanocolaou (1989). Emotion: A reconsideration of the somatic theory. New York: Gordon and Breach.

Andy addresses the fallacious study by Schachter and Singer (1962) that many believe confirmed the notion that emotion is a cognitive "labeling" of bodily states that are the same in all emotions. It isn't so. What is more, the old idea that if you shoot several people up with adrenaline they will all be in the same physiological state is equally fallacious.

Martin Taylor [920211], on behaviorism. What do you see in PCT that is of any value, if you think the flow of causality is as you described it: "The feedback loop involves a one-way information flow from intent to action to effect to perception." Who, among people who use PCT to model control by living systems, ever described such a path? Not one. Intent and perception meet in a comparison. Any difference between the two is "error," which drives action. And action alone does not detrmine the "effect" (whatever you mean by that) -- actions mingle with any other influences that affect the variable controlled by the control system. Actions plus disturbances determine the "effect," which is the reference level of the controlled variable.

The controlled variable influences perception, which of course meets the intention at the comparator. And all of this is continuous and simultaneous.

Find me a genuine behaviorist, of the variety that I earlier warned you are still alive and well, who describeds that loop going forward, backward, upside down, or any other way, and I will eat my wool hat!

I just don't get itDo you really think the control loop is just a linear causal path traced in reverse?

REQUEST FOR INFORMATION: Do any of the linguists, or any others for that matter, know of Victor H. Yngve, and his idea of "back-channel communication?" If so, in what regard are he and his ideas held and by whom?

Date: Wed Feb 12, 1992 3:35 am PST
Subject: Beer and Behavior

From Pat Williams (920212)

I'm disappointed! Especially with Rick M. and Bill P. Greg and I didn't start the conversation about Beer's bug, but since we knew a lot about it, thought a lot of it, and had discussed how it relates to PCT, we were glad to see it brought up. Greg warned me, but I was nevertheless surprised to see Beer's work being more or less dismissed as "cute" but not important, and basically "all wrong" although he had "by luck" incorporated some control into it. I was not only disappointed in the treatment of Beer's bug, but in the example it provides of how poorly some control theorists relate in general to the non-control-theory world. I really think that PCT is very important and should be much more widely known and utilized, but I sadly understand now at least one important reason why it is not.

I'm going to provide some history. When Greg came across Bill P. several years ago, I knew nothing about control theory beyond having read Wiener's book several years earlier. Greg was immediately very impressed, and since I find that Greg is almost always excellent in discerning what is really important, I took an interest too. Since that time, Greg has spent a lot of time and effort (for very little money or personal acclaim) promoting PCT by publishing and archiving existing documents. I'm a computer programmer (among other things). I do not claim to be a control theorist, but having discussed it in detail with Greg for years and having proofed all of the books and papers that have gone through Greg, I have a fair understanding. I still have a lot to learn, but I am basically very impressed with how far-reaching and elegant a solution control theory provides. I am not at all sure that all the details are correct (the evidence doesn't exist to prove or disprove them), but I'm convinced of the basic tenets.

When Greg came home with Beer's book, he was immediately excited by it (similar to how he was excited by Bill's work) and told me that we should implement it so that we could experiment with it. I basically dropped my major program that I was then working on (PictureThis), and implemented version 1 of NSCK in about two weeks. (Later versions only fixed some minor bugs and made it much faster.) I had been turned off by Artificial Intelligence (AI) about twenty years ago at MIT for two reasons: 1) The programs that some of the supposedly top AI people in the world were doing seemed based only on their verbal musings on how something might work (very little evidence or system) and were very kludgy and not at all robust because of this. 2) The main perpetrators of AI (at least at MIT) were obnoxious and arrogant and treated anyone who was not already in the "inner circle" as "knowing nothing." They wouldn't even let you take a course in the subject until you had already "proven" yourself. They acted very much like some PCTer's have been acting. I haven't followed AI closely since then, but from a distance, I haven't seen it improve much. Beer's work was different. It wasn't verbal, it was based to a large extent on what experimental evidence was available, it was robust, and (I admit) it was "cute." (I find myself rooting for the bug to find the food before it runs out of energy.)

The bug's "behaviors" can be talked about with words like wandering, edge-following, food-seeking, etc., but that is NOT how it is implemented. The bug just consists of a bunch of "neurons" connected together in simple ways. These neurons are simple analog devices (implemented digitally, of course) that have currents coming into them (some currents are negative), an internal voltage (which is set from the sum of the incoming currents and its present value), and a frequency output that is based on the internal voltage (it only varies if the voltage is above a certain threshhold). These frequencies are multiplied by some constant and fed as currents into other neurons. Sensor neurons also have currents coming into them from the "outside world" which indicate things like "smell" or "touch" in our terms, but are just more currents to the neurons. Motor neurons have outputs to the real world (forces to push the bug in some direction). There are also some neurons that generate random bursts of voltage. The neurons are based on real neurons, though they are not as complex. Each neuron might represent several neurons in a real bug (e.g., one neuron represents all the many that would represent an antenna touching something, but the effect is the same). These neurons operate in parallel (well not really, of course, on a serial computer, but effectively). It was quite fascinating to "debug" the parallel neurons while I was debugging my serial program. To REALLY understand (that word isn't strong enough, "grasp" maybe) how the bug works, you have watch how the various neurons' currents, voltages, and frequencies change at the same time (aided by the graphs on the side of the screen). You can talk about it in linear words, but it is not the same as REALLY grasping what is going on. (This same kind of "grasping" is needed to understand how multiple levels of hierarchical feedback loops are functioning.) I began to faintly grasp how the nervous system of a living being might really work. I also began to see how you could model all kinds of simple organisms using these neurons (and maybe a few different ones based on experimental data; we want to do this in NSCK 4.) NSCK 3 allows you to change all the constant values and connections for these neurons, add new neurons, and (if you can program) add more sensory and motor functions, so you can already do a lot. I was, and still am, very impressed with how well it works with just a few of these simple building blocks (neurons).

As I was implementing NSCK, I could see all kinds of negative feedback loops and even a hierarchical control system evolving. (I did it as Beer must have; first got individual neurons working, then put a few together to make a small circuit, then added more circuits, etc. I'll admit to not having read much of the book or getting an overview before I plunged in.) I started discussing with Greg how this obvious control was the same as and different from that of PCT. I wanted to know how the bug would have been implemented differently, using the same building blocks (neurons) from a PCT point of view. Some (but not most) of the bug seemed a bit kludgy and ad hoc, and the pseudo-physics bothered me. (Though, I'll admit it served its function, was much easier to implement, and made the program's execution faster. The bug robot is made on similar principles, and seems to work adequately with real-world physics, so I think my complaint is without too much basis.) We decided that the main difference between Beer's bug and PCT was the way lower levels of the control hierarchy were controlled: they were just inhibited instead of being passed a different reference signal. We concluded that it would be useful, interesting, informative, and do-able (but not at all EASY), to implement a PCT bug using NSCK. We would have liked to do it, but we just haven't had time (and I personally don't think I know PCT well enough). We also thought that Beer's bug might have been a bit better designed if he had had a PCT grounding to start with, but believe me, he might not have understood control in PCT terms, but he did UNDERSTAND it and UTILIZE most of the control theory tenets VERY WELL in practical ways (which is what was important to him and to most people that may have any use at all for control theory). It had nothing to do with luck.

Given that background, you can see why I am disappointed in some PCTers' treatment of Beer's work. I think it is important. I think it is an excellent example of how control theory can be used (forget the details). I think it would be an excellent way to try out the details of PCT on something that already works well with control, but isn't PCT-based. I think it (or similar models) would generate interest in PCT. (Beer's model has already generated a lot of interest. We have sold over 120 copies of NSCK to all kinds of different people from students wanting to use it for a school science project to entomologists to neuroscientists to science museums. It is being used in classrooms, etc. Note: we are not getting rich, we sell it for a $10.00 duplicating fee and give away the source code. We just did it for our own interest, and sell it because we want others to see how interesting it is.) There is one more reason that I think PCTers might be interested in Beer's bug. I think they might LEARN something, and not just how it does not follow PCT rules!

Now for my second main disappointment. I hadn't participated on the CSGnet before, though I had read some of what was coming across. When Beer's bug started to be discussed, Greg asked me to help him make replies since I knew more about it in some ways than he did. He wrote (with my input) all of the replies until now. But I did have a major input: when the discussion was getting a bit "nasty", I toned our answers way down. I wanted you guys to know that we thought PCT is great, but that others are doing good work too, and maybe if you analyzed their work in a POSITIVE manner, you might learn something, help others with THEIR work, and help PCT at the same time. I did not want to criticize you, because usually criticism just turns the criticized person "off", and he/she will not only no longer listen but tell others whatever you say is wrong without listening at all! You may be psychologists, but you do not understand everyday human psychology very well (or at least you do not practice your understanding). You frequently criticize FIRST and maybe LATER say "well, you had a little point right" when the criticized person is no longer listening. Telling anyone that all their past beliefs and actions are wrong is no way to win converts to PCT! Show them what they are doing right FIRST! (Everyone has something right!) Praise them for their interest in trying to understand how organisms work! Learn something from them! THEN if you still think you have something that could help THEM (why else would they be interested), tell them about it gently through examples, metaphors, and stories. If they "discover" it themselves it will mean a whole lot more than your trying to force it down their throats. (This is a very BIG reason that we homeschool our sons.) Use models as much as possible; they are much more effective than words, better than even pictures, and sometimes the only true way to understand something. I want PCT to be better known! You are acting like you don't!!

Pat Williams

From Greg Williams

>Rick Marken (920210) (Beer and Behavior)

>Beer's bug does produce cute behavior. But it is not based on useful
>underlying principles.

However, it is based, to a considerable degree, on empirical findings. And Bill Powers' hierarchical model for psychological control, with higher-level error signals resetting lower-level reference signals, is based on only the slenderest of empirical data. So here, I think is the impasse -- and a big part of the reason why PCTers have a difficult time getting published. To you and Bill, principles are of the essence -- they have to be, because they are virtually all you have, empirically speaking. A control loop does indeed offset disturbances -- but what about the myriad of details? The underlying notion of organismic control with internal reference signals is perfectly reasonable, as I've said before. I buy into that. But then I (and the referees) ask: where do you go from there? So far, you have made plausible models having only one or two levels, shown that those models can do certain things, and then acted as if you think that those things cannot be accomplished in any other way (or maybe "accidentally" by a non-PCTer). Beer's changing-loop-gain model also can do those sorts of things. And its changing-loop-gain technique is backed by considerable data (for various invertebrates). Why so much faith in the Powers model and its "well understood" principles, when there is no a priori reason to think that the actual control principles of nervous systems are likely to equal Bill's conjectures (or even that they are easy to understand)?

>Contrary to what Beer said to Gary Cziko, the bug (when a controlled variable
>is involved) is ALWAYS controlling.

The point was that the bug doesn't control in the continuous manner so beloved by PCTers who want to contrast continuous control with discrete stimulus and then response models.

>Maybe we could convince him that it might be worth it by building bugs that do
>what his do -- and then some -- using PCT.

I'll believe it when I see it.

>It would not really be hard.

I'll believe it when I see it.

Greg

Date: Wed Feb 12, 1992 4:42 am PST
Subject: points of view

[From: Bruce Nevin (Wed 920112 06:24:06)]

Me (920111 14:31):

>>In most cases, the basis for language control is in the spoken form, and
>>the visual aspects of the written record are secondary. Some confusion

Martin (920211 17:30):

>often expressed view. It is not at all clear, for example, that the original
>forms of written language had any connection whatever with the speech of
>the writers. Mostly the forms were related to visual impressions of concepts,

I see three points of view on these "how it is and how it got that way" questions, and a lot of confusion arising when we slip without notice from one to another:

1. How a living control system does it, in real time.

2. How a living control system learned it, acquired the skills.

3. How the conventions involved developed in the history of societies using them.

I intended point of view (1) and you responded in (3). Maybe the Etruscans (whence the Latin forms of the letters) still had some pictographic control of writing. Probably not. But the question and its speculative answer have little relevance to POV (1) for English today that I can see.

I agree that written forms of language and spoken forms tend to diverge. One reason is that written texts tend to be conservative, preserving features that have slipped out of currency in one spoken dialect or another. Another is that, by omission of detail, written forms mask (and perhaps retard) variation and change in language. An extreme is Chinese idiographic writing, where the same text may be "read" as very different sequences of words in the different socalled dialects (really, different albeit related languages). It's much more rebus-like than alphabetic or syllabary systems. But even there, as you know, some idiograms are used to indicate phonetic values, and that sort of usage was greatly extended in Japanese adaptations of Chinese writing; likewise for hieroglyphic writing.

The usage of what are normally pictographs for marking phonetic or phonemic indications, as indeed the very development of alphabetic and syllabary systems historically out of pictographic writing, actually attests to the essential primacy of the spoken form. And of course, from POV (2) it is always the spoken form that is learned first, and the written form that is learned on the basis of it. Helen Keller and perhaps others like her being all the more remarkably exceptional for that.

(Bill Powers (920211.0800) to Bob Clark 920210.0920) --

>evaluated judiciously. If you violate a social custom and everyone at the
>dinner table frowns at you, you might conclude that you've disturbed a
>social control system, which is exerting counterpressure. But if you
>carry the test to completion, looking for the means of control and the
>sensory channel, you will find only N individual control systems, each
>disturbed by your action. Blindfold all the people and they won't
>disapprove when you eat with your fingers. The social control system will
>disappear.

The logic here needs strengthening. To see that, apply the analogy on the basis of which claims of social control arise. Substitute elementary control system (ECS) within a living control system in place of person within a social system. Thus:

> If an ECS at level n "violates" a reference signal and ECSs at level n+1 issue error signals that return to it, you might conclude that you've disturbed a portion of a living control system, which is exerting counterpressure. But if you carry the test to completion, looking for the means of control and the sensory channel, you will find only N individual elementary control systems (ECSs), each disturbed by the action of the ECS at level n. Blindfold all the level n+1 ECSs--cut off the perceptual signal previously passed along from the ECS at level n--and they won't "disapprove" (issue an error signal) when the ECS at level n "violates" its reference signal. The living control system will disappear.

By blindfolding the other participants you have not proven that there is no control, only that the sensory channel of vision is part of the loop.

Now, what about that analogy?

It is possible, seems to me, to imagine each ECS to be a homunculus, a complete hierarchical living control system. We don't need all that power to model what an ECS does, all we need is three I/O channels, a comparator, an I/O function for each channel (but some of these are pretty sophisticated!), and a metabolic source of energy.

But notice that conversely we don't need all the complexity of hierarchical control to model a human being conforming to a social norm or convention. In conforming to a norm, a person seems to emulate an ECS controlling a perceptual signal.

Here I think is a key. People can *voluntarily* comport themselves as though they were ECSs in a social living control system. They act as if this is what they were. They do not do so all the time, they can participate in more than one such system, changing roles or carrying out more than one simultaneously through the ambiguity of what they do, and they do most of this without awareness of doing it.

Their goals of which they are conscious are idiosyncratic personal ones, the social roles, expectations, norms, conventions are merely tools at their disposal for attaining those goals. When you use a tool, you voluntarily restrict yourself to its limitations. When you drive a car, you constrain yourself to a seated position with relatively little freedom of choice about movements of your body. When you use a hammer you constrain yourself to hammering activities--and of course you constrain yourself to the nail-like possibilities in the world of perceptions, as Marcuse pointed out. When you use social norms so as to anticipate and exploit the conventionalized expectations of others, for the sake of cooperation in attaining your own personal goals, you constrain yourself to conforming in relevant respects to their conventional expectations. And you are likely to reify the social norms.

This equation of social role with voluntary service as a "social ECS" might be fruitful. What think you, sociologists?

Galloping in where angels fear, as usual. But "if a fool shall persist in rushing in . . . "

Bruce
bn@bbn.com

Date: Wed Feb 12, 1992 5:26 am PST
Subject: Two behaviorists reply

William M. Baum, SCIENCE 181, 1973, 1114 ff.

"Early behaviorists perhaps held views similar to the one Powers criticizes, but the inadequacy of describing behavior in terms of responses to stimuli was recognized over 30 years ago. With the recognition that behavior is affected by its consequences (the Law of Effect), open-loop descriptions began to pass away. Few behaviorists today would disagree with Powers's statement, 'there can be no nontrivial description of responses to stimuli that leaves out purposes.' Emphasis on purpose, in fact, has been the hallmark of modern behaviorists' thinking. The behaviorists' solution to the problem of purpose has been exactly the one suggested by Powers -- selection by consequences. That behavior and consequences constitute a feedback system is taken as a basic premise....

"Powers covers familiar ground in two other points. In his discussion of acts and results, he actually reinvents Skinner's concept of the operant.... Also familiar is the notion of the hierarchical organization of behavior.... Powers, albeit unwittingly, is square in the mainstream of modern behaviorists' thinking about instrumental behavior."

B.F. Skinner, SCIENCE AND HUMAN BEHAVIOR, 1953, Chapter 5:

"It is customary to refer to any movement of the organism as a 'response.' The word is borrowed from the field of reflex action and implies an act which, so to speak, answers a prior event -- the stimulus. But we may be able to make an event contingent upon behavior without identifying, or being able to identify, a prior stimulus. We did not alter the environment of the pigeon to ELICIT the upward movement of the head. It is probably impossible to show that any single stimulus invariably precedes this movement. Behavior of this sort may come under the control of stimuli, but the relation is not that of elicitation.... The unit of a predictive science [of behavior] is... not a response but a class of responses. The word 'operant' will be used to describe this class. The term emphasizes the fact that the behavior OPERATES upon the environment to generate consequences. The consequences define the properties with respect to which responses are called similar.... The OPERANT is defined by the property upon which reinforcement is contingent.

...

"The only way to tell whether or not a given event is reinforcing to a given organism under given conditions is to make a direct test. We observe the frequency of a selected response, then make an event contingent upon it and observe any change in frequency. If there is a change, we classify the event as reinforcing to the organism under the existing conditions.

...

"The connection between reinforcement and satiation must be sought in the process of evolution.... A biological explanation of reinforcing power is perhaps as far as we can go in saying why an event is reinforcing. Such an explanation is probably of little help in a functional analysis, for it does not provide us with any way of identifying a reinforcing stimulus as such before we have tested its reinforcing power upon a given organism."

...

"If we observe someone walking down the street, we may report this event in the language of physical science. If we then add that 'his purpose is to mail a letter,' have we said anything which was not included in our first report? Evidently so, since a man may walk down the street 'for many purposes' and in the same physical way in each case. But the distinction which needs to be made is not between instances of behavior; it is between the variables of which behavior is a function. Purpose is not a property of the behavior itself; it is a way of referring to controlling variables. If we make our report after we have seen our subject mail his letter and turn back, we attribute 'purpose' to him from the event which brought the behavior of walking down the street to an end. This event 'gives meaning' to his performance, not by amplifying a description of the behavior as such, but by indicating an independent variable of which it may have been a function. We cannot see his 'purpose' before seeing that he mails a letter, unless we have observed similar behavior and similar consequences before. Where we have done this, we use the term simply to predict that he will mail a letter upon this occasion."

[posted by Greg Williams]

Date: Wed Feb 12, 1992 7:00 am PST
Subject: All I Needed to Know in Life . . .

[From: Bruce Nevin (Wed 920112 09:18:54)]

. . . I learned in kindergarten, right?

(Joel Judd Tue, 11 Feb 1992 17:06) --

>If she cried in K[indergarten]
>class, the teacher comforted her; if she cried in Korean class, the teacher
>would become more "aloof" and tell her "the Korean equivalent of 'Don't be
>a baby.'" In another example the K teacher selected a Japanese boy to be
>the "leader" since his English was best. He began to translate and tutor
>other kids' English. Later when the Japanese teacher had to come in the
>class to discipline the Japanese group, she disciplined the boy in
>Saville-Troike's study because he was the oldest and therefore responsible

>for the group. According to Troike, this conflict engendered by the
>teachers alienated the two boys from each other for 3 weeks.

The same behavioral outputs constitute different "behaviors" in contexts defined by members of one culture or the other, and there is social agreement (conventionalization) in respect to each constitutive relationship (and in respect to cultural definitions of contexts) by those members, that is, by each of the people who use them. The conflict between two conventions had to have been experienced as conflict in the HPCT sense within each child. What is missing here is an account of how each boy attributed the occasion of internal conflict to the other boy. There is even less to go on to interpret the girl's crying. You say (Seville-Troike says) that it constituted "peer challenging" in the kindergarten class, but immaturity in the Korean class. What do you think her goal was in crying in the kindergarten class? Getting the teacher's nurturing attention for herself and away from others? Is that what S-T means by peer challenging, or something else?

>"   the criteria she
>used to judge relative readiness for first grade involved ONLY WHETHER
>RULES FOR APPROPRIATE BEHAVIOR HAD BEEN LEARNED AND WHETHER THE CHILDREN
>HAD DEVELOPED THE 'RIGHT ATTITUDES' ABOUT SCHOOL AND LEARNING. Whatever
>content knowledge was acquired during this socialization process was
>clearly a secondary consideration"

Academic success is constituted by conformity to the teacher's norms and expectations. That is, the teacher's rendition of what she thinks of as society's norms and expectations. The teacher functions as comparator. Her perceptions of them and her judgements of their behavioral outputs as conventional behaviors constitute the analog of perceptual input. Her control (or apparent control) of matters each child wants to control but may not be able to control or control well, and also to some degree her ability to help the child learn to control more competently, as rewards and sanctions, constitute an analog to the error signal. Children try to constitute relations to their environment (and especially to one another) on the analogy of the teacher's relationship to them. (Of course, this is analogous to their parents' relationships to them.)

This continues long after kindergarten. I just last week got a letter from the graduate chair at Penn. He will be de facto chair of my dissertation committee. In essence, he said that I would have to be more subservient and agree that my teachers' views were correct, or he would not sign my dissertation. Fortunately, I regard the dissertation as an academic exercise with no great existential import for my identity. I made a most politic reply.

>To the extent that this teacher (or any similar teacher) can effectively
>judge such criteria, and Saville-Troike argues that she did (based on cases
>with other students), it would be interesting to know HOW she did it, if
>she was using some sort of "folk" knowledge of the Test. This also implies
>that the teacher is complicit in furthering the "right attitudes" fostered
>by those in charge of the school, etc.; that she agrees with such
>perceptions of schooling and has some form of them herself by which to
>judge student behavior.

She has internally-maintained reference signals for conformity to social norms. Many of these I believe constitute models of how individuals in dyads, small groups, and large groups ought to comport themselves. Just as she is able to identify a role for herself in many situations out of one of these models, as applied to the situation, in the same way she is able to see how children in a situation fail to behave as predicted by the model. By teaching appropriate behavior with reference to the other participants and with implicit reference to the model, she teaches all the children the model. They learn its generative applicability to novel situations, they learn to assume various of the roles that it affords, they learn to preferentially choose some roles over others (to compete, to be the quiet observer, to fail, etc.) And so it goes.

>Who wants to bother with the intricacies of another language when you've
>already got one, and the second is "good enough"?

So what would make their restricted competence in L2 no longer "good enough"?

-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-

(Tom Bourbon Wed, 12 Feb 1992 00:07:22 -0600) --

>   REQUEST FOR INFORMATION: Do any of the linguists, or any others for that
>matter, know of Victor H. Yngve, and his idea of "back-channel communication?"
>If so, in what regard are he and his ideas held and by whom?

I know the name, but not much of the work. Emeritus, I think. In the sociology of the field, I think marginal, though that of course says more about the field than about the intrinsic value of his ideas. Is this "back-channel" idea just another take on nonverbal communication? Reference?

Date: Wed Feb 12, 1992 7:18 am PST
Subject: Re: Behavior

[Martin Taylor 920212 10:00 Early in the morning for me!]
(Rick Marken 920211)

Rick thinks he didn't make himself clear. I think it is I who failed in clarity because I have no problem with any of the positive things Rick said or says. But Rick's response is so far away from the intent of my posting that I hardly know where to begin in clarification. Very briefly (since I am overdue in a meeting), I intended no claim about what other psychologists believe. My intent was to argue that IF the PCT model is correct, then a behaviourist approach will arrive at the same structure, because the place where stability exists is in the relation among the percept, the reference, and the error signal in an ECS. The PCTer asks "what does the person perceive, in relation to what he wants to perceive" and the behaviourist asks "what does the person do in respect of what he wants to happen." The PCTer takes the error signal as the derived variable, the behaviourist takes the percept as the derived variable. Rick has made this plain some months ago, and I am surprised that my attempt to put his argument in other words generated such a strong antipathetic response.

More later, after my meeting--perhaps.

Martin

Date: Wed Feb 12, 1992 9:13 am PST
Subject: Yngve

I received a catalog today from Indiana University Press, and in it I see a book by Yngve. Here is the info in the catalog:

Yngve, Victor H. Linguistics as a science. Regular 9.95, on sale
    for 2.95.
    ". . . a very interesting, stimulating and thought-provoking
    book . . ." --Journal of Phonetics.

Apparently that was the most impressive supporting quote they could come up with. Is that where the back-channel business comes up?

Another book on sale there:

Marjorie Harness Goodwin. He-said-she-said: talk as social organization
    among black children. 19.95 on sale for 13.95.
    ". . . one of the best, if not the best source, for anyone interested
    in how boys and girls use language in their daily lives. . . ."
    --Deborah Tannen, author of _You just don't understand: women
    and men in conversation_

Indiana U. has a tradition of supporting people and lines of thought perceived as being out of the mainstream.

        Bruce
        bn@bbn.com

Date: Wed Feb 12, 1992 9:46 am PST
Subject: interesting language disability

[From: Bruce Nevin (Wed 920112 12:00:57)]

Cross-posted from
Linguist List: Vol-3-141. Tue 11 Feb 1992.
Posted there by: rws@mbeya.research.att.com (Richard Sproat)
-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-

Gene Controls Learning Of Grammar, Researchers Say
By PAUL RECER
AP Science Writer

CHICAGO (AP) - A single dominant gene controls the ability to learn grammar, said a researcher who studied a family whose members don't know to add ``ed'' for the past tense of verbs or ``s'' for plural nouns.

Myrna Gopnik of Montreal's McGill University said Monday the studies show that in all other ways the members are intellectually normal.

But, she said, ``Language is a problem they solve by sheer wit.''

She said people lacking the grammar gene ``are worn out just by talking'' because they must continually struggle with verb tense and noun plurals.

``The hardest part for them is people thinking that they are stupid,'' Gopnik said. ``They are not. You have to think of them as people without a native language.''

Gopnik reported on her research at the meeting of the American Association for the Advancement of Science.

Steven Pinker of the Massachusetts Institute of Technology said it was first suggested in the 1960s that there was a genetic component to learning language and that recent studies, such as that by Gopnik, supports that belief.

He said his research shows that learning words and learning to apply the rules of grammar are two different functions springing from different parts of the brain.

People with a normal grammar gene naturally learn language rules that make verbs past tense or turn nouns into plurals, Gopnik said.

People lacking the gene, however, must learn through another intellectual process how to change verbs and nouns.

``They will say `today I walk, yesterday I ...' and they don't know how to finish,'' said Gopnik. ``For some reason they don't build the general rules of language'' such as adding the ``ed'' to mak walk past tense.

Gopnik studied a family in which, of 30 members across three generations, 16 suffered from the inability to learn grammar rules. Otherwise, the 16 were normal, as were the unaffected family members.

``They have not found a genetic marker for this gene,'' said Gopnik. ``But we have very good evidence that it is genetic.''

Inheriting the language defect follows classic rules of genetics, said Gopnik. A member of the family has a 30 percent chance of inheriting the language problem, while the rest of the population has a 3 percent chance.

Studies by other researchers have supported the existence of a grammar gene, she said.

One showed that if one identical twin suffered a language difficulty, there was an 80 percent chance that the other would also, said Gopnik. For non-identical twins, there was only a 35 percent chance that both would have the problem.

Thomas G. Bever of the University of Rochester said his research shows that families with left-handed members may inherit a tendency to process language differently.

He said that righthanders with no left-handed family members are more senstive to the grammatical structures of language than are righthanders with left-handed relatives.

Right handers with some left-handed relatives can recall and use individual words more skillfully than members of an all right-handed family, said Bever.

This suggests, he said, that there is an inherited or genetic component to naturally learning grammar rules or in the ability to use individual words.

AP-NY-02-10-92 1749EST

-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-

I wonder if other reductions are affected, such as putting -ly on a nonsense adjective to make it an adverb (her running was gelfan and swift, and as she ran gelfanly over the field, etc.). These things are not going to be evident without active search for them, since the affected people will have found many ways to circumvent their lack and still communicate effectively--like the students whose limited L2 command is adequate to get by. Or like the "mentally retarded" person who can't tell time, but finds out the time at need by the clever expedient of wearing a watch that is broken: It provides a foolproof excuse for asking someone what time it is. Appearing normal and unexceptional to others seems to be the controlled perception

Bruce
bn@bbn.com

Date: Wed Feb 12, 1992 12:43 pm PST
Subject: Re: emotion; behaviorism. Yngve

[Martin Taylor 920212 14:00]
(Tom Bourbon 920212 01:00)

I really must have been a model of unclarity in my "behaviourism" response to Rick. Tom Bourbon now seems to think I don't know the first thing about control systems and the PCT view. On the contrary, I feel that for the last six months I have understood it quite well, and that recently I have begun to understand why PCTers have such objection to classical behaviourists, as well as why that objection does not apply IN PRINCIPLE to behaviouristic approaches. I tried to introduce that clarification yesterday, but my writing seems to have been as clear as mud. I'll try again.

Yes, Tom, you warned me that dinosaurs exist, of the variety that I had been told 35 years ago was long extinct. I do not mean coelacanths when I talk of fish.

You say:
> I just don't get itDo you really think the control loop is just
>a linear causal path traced in reverse?

No, of course I don't. That would be stupid, wouldn't it? I may make mistakes, and I hope I do, because otherwise I won't progress. But that kind of notion is wildly far from my way of thinking.

Rick came closer when he (thinking he was cricitizing me) asked something like "But what is that people really DO?" That was the question I was addressing, and arguing that the answer is accessible by tracing the anastomatic net of percept and reference in either direction. I started from the proposition that the PCT model was correct, and asked what approaches were feasible to study it, taking as given that the ONLY experimental data were those observable by third parties, not introspection. If the PCT model is wrong, then my posting of yesterday is quite wrong. It works ONLY within the context of PCT.

I hope this helps to clarify yesterday's posting. If not (and maybe if so), I will try to get at the issue another way.

Martin

Date: Wed Feb 12, 1992 12:55 pm PST
Subject: Re: points of view

[Martin Taylor 920212 14:30]

(Bruce Nevin Wed 920112 06:24)

Bruce, I don't think CSG is (at the moment) the place to get into a side discussion of the relation between spoken and written language. We must agree to disagree--and I am sure you realize, as I do, that your description of the use of Chinese characters for phonetic purposes is as oversimplified as mine. I agree that all your points 1, 2, and 3 are valid points of discussion. I think the discussions interact, and the evidence you adduce to support the notion that written language is more or less a means of transcribing spoken language is much the same as the evidence I would use to say that it is not. The interaction is much more subtle. If you want to discuss this further, let's do it off-line.

I am amused to see you take the position in regard to social control for which you castigated me last year. Perhaps progress is possible.

Regards
Martin

Date: Wed Feb 12, 1992 1:00 pm PST
Subject: DEMO 2 NOW AVAILABLE!

[from Gary Cziko 920211.2200]

I am pleased to announce that Bill Powers's Demo 2 ("Introduction to Control Theory Part 2: The Theory of Control") is now available electronically, thanks to the cooperation and assistance of Bill Silvert of the Bedford Institute of Oceanography in Dartmouth, Nova Scotia.

Demo 2 (as Demo 1 previously announced) will run on IBM PCs and compatible machines with 286 (AT class) or "better" CPUs and a mouse. Demo 2 will demonstrate the operation of a perceptual control system and allow the user to fine tune the model to fit data from a human subject. It is organized in a step-by-step tutorial format with explanation alternating with hands-on demonstrations and data analysis.

Rick Marken (920120) has observed that . . . sorry, we don't need to hear all that again, do we?

There are two ways to obtain Demo 2:

First Method: File Transfer Protocol (ftp)

Using ftp is the easier of the two ways to obtain the program as you will receive a single binary file ready to run on your PC. (My understanding is that just about everyone on Internet can ftp. Bitnet and MCI Mail users may have to use e-mail as described below.) If you know how to obtain files via anonymous FTP, all you need to know is that you want the binary file dem1a.exe on the pub/csg subdirectory of biome.bio.ns.ca.

If you have not used ftp before, here is a guide to getting this file:

Start off by typing:
ftp biome.bio.ns.ca (or ftp 142.2.20.2)
You will get a login prompt. Login as anonymous.
You will get a password prompt. Enter your last name (this is just a courtesy).

You will then receive a message, followed by a prompt. Issue the following commands:

```
cd csg
binary
get dem2a.exe   (This copies the file to your machine)
quit
```

After you get the file dem2a.exe off your mainframe, put this file on a subdirectory on your hard disk or on its own otherwise empty floppy and type "dem2a.exe". This will "explode" the zipped file into a dozen or so program files. Then to run Demo 2 type "demo2" (but make sure you have your mouse connected and active first). You can then discard dem2a.exe (or better yet, archive it) since it is no longer necessary to run Demo 2.

You can do the above directly from your PC if it is connected to Internet in which case the file will be transferred to your machine's active disk. If you are using a terminal connected to a "mainframe" machine, you then need to figure out how to get it off your mainframe onto a PC diskette readable by your computer (your local computer services office should be able to help you with this, but make sure they know that it is a BINARY file).

Second Method: Electronic Mail (e-mail)

Anyone reading this message (including those using MCI Mail) should be able to obtain Demo 1 via e-mail. You cannot send binary program files via e-mail since e-mail systems expect to receive regular ASCII ("typewriter") characters and binary files contain non-ASCII information.   Therefore, the file dem2a.exe has to be encoded into ASCII before it can be sent. The program used to change program (binary) files to mailable ASCII files is called uuencode and the program used to turn the ASCII encoded files back into binary programs files is called uudecode.

So, to get the uuencoded version of Demo 2, you just send a ordinary mail message to server@biome.bio.ns.ca and include in the message the line:

uuencode csg/dem2a.exe

You will then receive probably four files from the mail server. You need to concatenate these files in the proper order and decode them. If your system has only the standard uudecode program, then you have to edit the concatenated file to remove the headers, so the file starts with a line that says 'begin 644 dem1a.exe' and ends with a line saying just 'end', with a whole bunch of cryptic lines in between (most of these lines begin with M, but there are two other lines that you have to leave in just before the end statement). You can use a word processor for this editing and concatenation, but make sure that you save the file in ASCII format (also sometimes called nondocument or plain text format). However there are many enhanced decoders that can do this automatically, so you just have to save the parts in order into a file and decode them without editing.

* *

We plan to also make Bill Powers's Arm demonstration ("Little Man") available when the new version is ready. While it may seem like a hassle to get these files, this is in the long run much more convenient than mailing floppies. Anybody taking the time to read CSGnet traffic should definitely work through Demo 1 and Demo 2 to get a feeling for what perceptual control is and what the basic PCT model is.

Finally, I attach Bill Powers's invoice form for these programs. Since Bill is now living on fixed retirement income, I'm sure he would appreciate being properly compensated for his work on these marvelous programs. If only half a million of these programs are purchased, Bill will be able to buy an old hotel in downtown Durango for conversion to the much needed Institute for Control Systems Research!

==

INVOICE

Invoice number:_______________ (use YYMMDD)

Date:_________________________

Remit to: William T. Powers
73 Ridge Place, CR 510
Durango, CO 81301

To: _________________________ (purchaser)

____

____

____

Conditions: This invoice covers the purchase of any of three programs, named Demo1, Demo2, and Arm, which are in your possession on approval. Payment for use of these programs by individuals for their own enlightenment is optional. For institutional or professional use, the charges (taxes included) are as follows:

Demo1: The phenomenon of control
$35 per class-semester (or course) or other professional use
$150 for unlimited use in teaching by a single department

Demo2: The theory of control
$60 per class-semester (or course) or other professional use
$150 for unlimited use in teaching by a single department.

Arm: A control-system model of pointing behavior, version 1
$35 per class-semester (or course) or other professional use
$150 for unlimited use in teaching by a single department.

Copies of these programs may be made freely for student use or for any other noncommercial and nonprofessional use. The programs must be distributed as received with no changes, and may not be sold.

Program	Unit price	# courses	Total remitted
Demo1	$35	______	__________
	$150	(unlimited use)	__________
Demo2	$60	______	__________

	$150	(unlimited use)	__________
Arm	$35	______	__________
	$150	(unlimited use)	__________
		Grand total	__________

Thank you for your order.

Gary A. Cziko   Telephone: (217) 333-4382
Educational Psychology   FAX: (217) 244-40538
University of Illinois   Internet: g-cziko@uiuc.edu
1310 S. Sixth Street   Radio: N9MJZ
210 Education Building
Champaign, Illinois 61820-6990
USA

Date: Wed Feb 12, 1992 1:07 pm PST
Subject: ATC--An opportunity?

[Martin Taylor 920212 14:40]

This morning I attended a discussion led by one of our people who has an interesting problem of modelling, one that I think might serve as an opportunity for PCT and the CSG members. This is a request for comments and suggestions.

The situation is a small Air Traffic Control (ATC) game on the Macintosh (not a commercial one, but one programmed locally). The screen display consists of a circle with diagonals drawn N-S, E-W, NE-SW, and NW-SE. There are thus 8 points where the diagonals meet the circle. The circle represents the air space controlled by the subject of the experiment.

In addition to the airspace display, there is a table of aircraft descriptors, and a rosette that is used to set the heading of any aircraft and a slider control to set the aircraft's altitude. These latter two are used by clicking on them with a mouse. They affect an aircraft that has been selected from the airspace display by clicking on it with the mouse.

In the experiment, an aircraft is shown as a little arrow that points in the direction it is headed, which may be one of the 8 major directions. The aircraft enters the airspace either at one of the 8 cardinal points at the circle boundary or by taking off from one of the airports located in the "controlled" airspace. Each aircraft has an associated destination. It must leave the airspace at a specified (cardinal) point at a specified altitude, or it must land at a specified airport. So long as the subject does nothing, an aircraft maintains its heading and altitude. The subject can click on an aircraft icon and then on the heading rosette or the altitude slider. The aircraft will change course by no more than one 45-degree increment in each screen update period (which is changed as a factor in the experiment between 3 and 15 seconds in different runs).

The subject's task is to get each aircraft to its destination by a minimum distance path, without crashing into any other aircraft.

The discussion centred on how to model the operator, using a task network modelling language called microSaint. It occurred to me that this situation is a natural one to model as a control hierarchy, with loops for mouse movement, for various perceptions relating to the aircraft and their relations to one another and to their target destinations, and so forth. But what would one disturb in order to tease out which percepts are really being controlled? The aircraft movements are ballistic, and not under continuous control. Making them deviate from their ballistic path would alter the ground rules of the environment, and if their tracks turned out to be controlled under such disturbances, we would gain no information about what was being controlled under the ballistic regime.

The experimental results so far are that people make few errors and are not far wrong on the minimum distance routes, so long as the number of aircraft on the screen remains below some critical value. Above that, errors mount and it seems as though the number of well-controlled (?) aircraft remains constant. So we are dealing with a condition of resource limitation, which is common in the real world. (Come to think of it, I don't remember any CSG discussion of resource limitation in the last year. Has there been any?)

It seems reasonable to me that we are dealing with a question of stability in a hierarchic control system. If we differentially delay, say, the response of the aircraft to a click on the heading rosette, or the time for which a selection of an aircraft remains valid, or things like that, we should probably be able to affect the stability and bandwidth of lower loops, and this shoudl affect higher loops. But that's not disturbance as usually discussed in connection with PCT. There may be tracking going on, but for the most part it is in the imagination loop somewhere, until a plane deviates too far from its desired track AND is noticed to do so.

Do CSG people see an opportunity here? The experimental source code is available, effort is being placed on simulation modelling at a quite detailed level (if data justify the detail), and good ideas for teasing apart different possible models of the subject would be most welcome. In particular, I would like to segregate models that depend on sequential queues as the location of the resource limitation from those that depend on the characteristics of control systems (S-R vs PCT, to be crude).

By the way, I have nothing to do with this project, except insofar as one of my prerogatives as "Senior Experimental Psychologist" is to interfere with what other people are doing.

Martin

Date: Wed Feb 12, 1992 1:19 pm PST
Subject: Beer, behavior and reconciliation

[From Rick Marken (920212)]

Well, it looks like I really created some flack this time.

First, let me thank Tom Bourbon (920212) for affirming the fact that I am not the only one who had trouble understanding Martin Taylor's [920211] post on behaviorism. I concur completely with Tom's reply to Martin. Maybe Martin could explain his position in a little more detail. Maybe he just means that if behaviorists understood PCT they would know that behavior is controlled variables.

Pat Williams (920212) says:

>I'm disappointed! Especially with Rick M. and Bill P.

> Greg warned me, but I was nevertheless surprised to see Beer's
>work being more or less dismissed as "cute" but not important, and basically
>"all wrong" although he had "by luck" incorporated some control into it. I was
>not only disappointed in the treatment of Beer's bug, but in the example it
>provides of how poorly some control theorists relate in general to the non-
>control-theory world. I really think that PCT is very important and should be
>much more widely known and utilized, but I sadly understand now at least one
>important reason why it is not.

I was afraid this might happen. Greg had hinted that we might be able to encourage Beer toward the PCT fold by helping him understand how close his models were to PCT models. Gary Cziko suggested posting the "Beer bug" thread to Beer -- without Marken's peurile (my word) digs. I have heard Pat's complaint ("PCT would be better known and more widely accepted if PCTers were nicer to the 'opposition'; empasize the similarities, play down the offensive differences"). The idea, of course, is to encourage people, especially talented people, who are "close" to PCT, into the fold.

I am willing to admit that I can be a kind of contentious sort. Maybe I should write more sweetly and friendly. I accept my scolding from Pat and Gary and whomever else (and there have been others) -- I'll try to be nicer. Maybe. But I strongly reject the idea that, by being more conciliatory we can encourage more people toward PCT. It won't work! The people who would "accept" PCT as a result of verbal encouragement don't really get it anyway. These people are like religious converts -- and their "acceptance" of PCT is as fragile as that of any religious convert. If another prophet comes along with a better story then off they go.

The "real" converts to PCT are the one's who have been convinced by their own observations and reasonings. They have had what I call the "Holy s**t" experience. They have realized that behavior is the control of perception; they know what that means experientially, empirically and theoretically. Once you've had the "Holy s**t" experience, you are a PCTer; and there is no going back. There is no way to make a person have this experience (as any PCTee would know from their own theory).

The point is that a real PCTer becomes a real PCTer ON THEIR OWN. No one can (or should) convince you that PCT is right. Only nature and reasoning can do that. I became a PCTer on my own; I didn't get in touch with Bill Powers until I had done enough studies and modelling to convince myself that I understood the basic insight of PCT and that I knew how to demonstrate the basic phenomena. I had already, on my own, had the "Holy s**t" experience.

Bill Powers is a wise fellow. In his response to my first, enthusiastic letter to him, he said "getting into PCT can be a pain in the ass. Don't do it unless you REALLY know what you are getting into". He tried to DISUADE me from getting involved in PCT.

Now I understand why. We don't need converts -- they tend not to understand what they are doing. Converts are people like Carver and Schier, Lord, and so on. These people have not had the "Holy s**t" experience. Once you get into PCT, you can tell who has had it and who hasn't; the one's who haven't don't apply PCT properly and eventually move on to conventional s-r

behavioral science -- perhaps using the language of PCT (Glasser comes
to mind).

People who have had the "Holy s**t" experience will not be turned off to
PCT by mouth-offs like me -- because they know the value of the model.
And people who are going to have the "Holy s**t" experience will not be
turned if somneone like me says that their current ideas are WRONG. If
they don't already sense that, they won't process PCT anyway. If they do
sense it there response will not be defensive-- it will be "how is it
WRONG?" and then evaluate what I say and take it or leave it. Similarly,
all the sweet talk in the world will not produce a "Holy s**t" experience
in someone who is not interested in understanding PCT.

Pat suggests that PCTers are being arrogant (like some AI types) by not
seeing the value in work that has not been done from a PCT perspective.
I am sorry if that is how it seemed. I do not feel arrogant; I don't
mean to dismiss Beer's work. But there's LOTS of clever work being done
out there -- expert systems, robotic systems, etc etc. I have no doubt that
Beer and Brooks and all the others I have discussed are orders of magnitude
smarter than I am. But they don't know PCT.

>The bug's "behaviors" can be talked about with words like wandering, edge-
>following, food-seeking, etc., but that is NOT how it is implemented. The bug
>just consists of a bunch of "neurons" connected together in simple ways.

I know-- that was not my point.

>As I was implementing NSCK, I could see all kinds of negative feedback loops
>and even a hierarchical control system evolving.

I know.

>   We decided that the main
>difference between Beer's bug and PCT was the way lower levels of the control
>hierarchy were controlled: they were just inhibited instead of being passed a
>different reference signal.

I think there is NO difference between Beer's bug and PCT - Beer's bug is a
set of componenets that can be used to implement various perceptual control
systems. The problem I have with it is only that the architecture makes it
hard to see what perceptions are controlled, and how. There is nothing
wrong with the bug -- and your program simulation of it is great. It's how
Beer thinks about what he is doing that, I think, in the end will limit
his ability to build interesting robots.

>   We concluded that it would be useful, interesting,
>informative, and do-able (but not at all EASY), to implement a PCT bug using
>NSCK. We would have liked to do it, but we just haven't had time (and I
>personally don't think I know PCT well enough).

Sure you do. We could use NSCK as is to build bugs based on PCT. It would
just be a bit limited; because the bug doesn't have many perceptual functions.
I think it would be better to start from scratch and buiild a PCT construction
set that could give the bug all kinds of perceptual functions -- that would
work in the simple simulated world of the bug.

>We also thought that Beer's
>bug might have been a bit better designed if he had had a PCT grounding to
>start with,

If his components are really like neural components then they are just fine -- a PCT grounded construction set might be inclined to make the componenets more functionally oriented; giving up fidelity to the neural facts in order to make the architecture easier to understand in terms of control system componenets. That's why I suggested a compiler -- from PCT components to Beer bug componenets.

> but believe me, he might not have understood control in PCT terms,
>but he did UNDERSTAND it and UTILIZE most of the control theory tenets VERY
>WELL in practical ways (which is what was important to him and to most people
>that may have any use at all for control theory). It had nothing to do with
>luck.

Luck may be a poor term. But I don't think he understands why his control systems work.

>Given that background, you can see why I am disappointed in some PCTers'
>treatment of Beer's work. I think it is important.

Why is it more important than the s-r machines (which are really control systems) that are descibed in the article on flocking birds that I mentioned. Many people have built clever robots that control variables. I just don't count it as significant if they don't understand the basic organizing principle -- control of input. Without that explicit conception you get side-tracked. That's why AI and robotics have been going on for so long with only haphazard success -- but success none the less. I am impressed by all these technical achievments. But I think these achievements could be far more formidable if people understood the underlying principle of control.

>There is one more reason that I think PCTers might be interested in Beer's
>bug. I think they might LEARN something, and not just how it does not follow
>PCT rules!

It's not a matter of following PCT rules or dogma. To the extent that it controls then it follows PCT "rules". But the architecture makes it hard to see how it follows those rules.

>   Telling anyone that all their past beliefs and actions
>are wrong is no way to win converts to PCT!

You bet. I don't think we do that. But I don't think we want converts -- I sure don't. Note the first part of this post.

>   Show them what they are doing
>right FIRST! (Everyone has something right!)

Everyone (even my mother) is right about some things -- no question. I think we do try to point out what people have right. But they are not going to be interested in learning PCT anyway unless they themselves are experiencing some problems. If they think they are right about what we think they are wrong about then there is no chance that we will convert them -- why should we want to anyway. If they are happy with what they are doing, I want to keep them happy (as long as it doesn't hurt anyone else -- Hitler was happy with his ideas but those ideas led to him to hurt other people so he should have been converted --or shot). Most psychologists are happy with what they are doing -- they rarely hurt people (I think) and I have no interest in converting them. I just want to put my ideas out there so that the 1 person in 10000 who might be able to get into PCT can do so.

> If they "discover" it themselves it will mean a whole lot more than
>your trying to force it down their throats.

It will be the only thing that means ANYTHING.

> I want PCT to be better known! You are acting like you don't!!

If PCT is anything, it is a revolutionary new way of looking at behavior. We could try to conceal that and make it seem like it's compatible with a lot of conventional wisdom--but it's not. Once you start explaining PCT to people who hold dear one or another conventional concept that conflicts with PCT, you start sounding like a contrarian -- unless you lie. I want PCT to be know, understood and used. But I won't say that PCT is really compatible with this and that concept, when it's not.

Beer describes his bug in behaviorist terms. He wires it up with neural connections that are hardly revolutionary and ends up, as a side effect, producing systems that control some perceptual variables -- and he describes this in s-r terms. What am I supposed to do???? Say how much I like what he's got -- up to the point where he misses the fact that he's building perceptual control systems? After I tell him all the nice things about his bugs, when do I tell him that he's missed a point. How do I say this without implying that I'm right and he's wrong. There's just no way out; the best (as Bill has said) that we can hope for are constructive debates. But if people are offended when you say that they are wrong and explain why, what can I do -- other than say they are not wrong (which, from my perspective would be a lie). I don't mind when people say PCT is wrong. In fact, I wish people would do it more. Explain and demonstrate what is wrong with PCT. Some people have but not many do. I take criticism of PCT as an opportunity, not as a threat. If Beer has great ideas then he should take my critiques as an opportunity to explain how I am wrong. Again, this ain't a religion, folks.

>From Greg Williams

>>Rick Marken (920210) (Beer and Behavior)

>>Beer's bug does produce cute behavior. But it is not based on useful
>>underlying principles.

>referees) ask: where do you go from there? So far, you have made plausible
>models having only one or two levels, shown that those models can do certain
>things,

What more has Beer done?

> and then acted as if you think that those things cannot be
>accomplished in any other way (or maybe "accidentally" by a non-PCTer).

This is ambiguous. If the system controls, then, yes, it can only be accomplished by a control organization (that's the WAY). But the functions of the control organization can be implemented in many WAYS -- wires, tubes, neurons,etc.

> Beer's
>changing-loop-gain model also can do those sorts of things. And its changing-
>loop-gain technique is backed by considerable data (for various
>invertebrates).

No problem -- it's just that it is a control organization -- and, to the extent that it controls it controls PERCEPTION. Beer has missed that point. It's important because what a control system DOES depemds on what it PERCEIVES.

>Why so much faith in the Powers model and its "well
>understood" principles, when there is no a priori reason to think that the
>actual control principles of nervous systems are likely to equal Bill's
>conjectures (or even that they are easy to understand)?

It's that ONE principle that's missing -- what is perceived determines what is controlled. I'm not against the implementation. I don't mind implementing control by loop gain adjustment, digitally or analogically or whatever. I just think that Beer doesn't know that he is building a device that controls its perceptual inputs. If he knows that, great. It just didn't seem like it from the article and book.

If you are saying, however, that what Beer calls "adaptive behavior"(which I take to mean the production of consistent results in a disturbance prone environment -- ie control) can be implemented in a way that DOES NOT involve the control of perceptual variables, then we have a HEAVY difference with Beer. I'm saying that Beer is building machines that control -- he's just not using our architecture and he is not recognizing that these machines control inputs. I think you are saying that Beer is building machines that control and the fact that he doesn't know that they control inputs is not important. OK. But if you are saying that Beer's machines control without controlling perception than I want to see how that can work.

Best regards (really)

Rick

Date: Wed Feb 12, 1992 3:47 pm PST
Subject: Scolding;behaviorism

[From Bill Powers (920212.1200)]

Many important messages in today's posts, perhaps the most important one being Pat Williams' scolding. I feel like defending myself, which is a sure sign that she has hit the mark.

She says "Telling anyone that all their past beliefs and actions are wrong is no way to win converts to PCT!" This is correct. I think that the essence of Pat's complaint is that we indulge in "bashing." And we do. If we continue doing this we will end up talking to each other and nobody else.

There's nothing wrong with writing hate mail to behaviorists and others who give us a hard time. But once it's written, we should delete it. Such attacks reflect both frustration and vulnerability -- frustration at being treated, often, just as Pat says we treat others, and vulnerability in knowing that we have not gone very far toward building a real CT science of human and animal nature.

We don't have to refrain from criticizing ideas we think are wrong. But I think we have to be most careful to limit criticism to specific points and steer clear of anything that sounds ad hominum. Even when we want to

grab the guy by the neck and throttle him.

Martin Taylor (920211.1745) --

Rick has already responded in his passionate way to this, but I need a turn, too, even if my comments are largely redundant with Rick's:

>I take for granted that the data of any psychology must be the
>observation of what people (or animals) actually do.

As Rick said, the basic question is what you're observing when you say you're observing what people do. The basic concept of control theory is that patterns of behavior regular enough to recognize are not simply "emitted," but result from an active process of controlling for just those patterns. When you reach for a glass of water, there is no "reaching" signal entering the muscles. In fact, the main forces generated by the signals entering the muscles are aimed upward, against gravity. These forces fluctuate as clothing drags on the arm or, if you're on a cruise ship, as the magnitude and relative direction of the gravity vector varies. These extraneous forces, however, don't materially affect the reaching -- the pattern of behavior that we see remains essentially normal. But the action of the nervous system on the muscles is continuously varying and may not ever repeat.

You say

>The feedback loop involves a one-way information flow from intent to
>action to effect to perception (no, I'm not forgetting it is
>continuous). The observable things are in the range action to effect.
>I don't see any obvious reason why there should be a great benefit to
>following the flow forward or backward, if both routes arrive at the
>same place.

Not quite -- the intent is not converted into action. The error is, and the error depends on both the intent and the current perception.

The "obvious reason" is that you can't follow the flow backward, except under very restricted and artificial circumstances (which are created, as nearly as possible, in behavioristic experiments). At every step in the following chain, disturbances can enter that are independent of what the nervous system is doing:

```
  NERVOUS SYSTEM -> MUSCLE TENSION->JOINT ANGLE->APPLIED FORCE->POSITION
                  /                 /            /              /
                 D                 D            D              D
```

The regularities we observe as behavior occur at the END of this chain, not the beginning. If you follow the chain even further outward, to "behaviors" such as steering a car, mixing a cake, walking in a figure eight, or ice-skating, even more disturbances come into the picture and tracing the chain backward becomes even more obviously impossible. In every example, the greatest regularity is found at the end of the chain.

One last point, to clarify Rick's answer to this:

>PCT people believe that what people do is done because that will allow
>the people to set up desired perceptual conditions.

In fact the doing itself is one of the perceptual conditions that is produce by action. It's easy to forget that this is a hierarchical system, with no action other than production of muscle contraction occurring open loop. We don't really see "actions" when we observe behavior. Mostly, we see controlled consequences of action.

In a later post:

>My intent was to argue that IF the PCT model is correct, then a
>behaviourist approach will arrive at the same structure, because the
>place where stability exists is in the relation among the percept, the
>reference, and the error signal in an ECS.

But what about the stability of the controlled variable against disturbances? That, too, has to be explained. Perhaps you should lay out just what the behaviorist approach to explaining behavior is. There's got to be some difference between the behaviorist approach and the CT approach, or the behaviorist approach would have uncovered the phenomena of control long ago.

Greg Williams (920212.0620) --
RE: Two behaviorists.

Baum:
>The behaviorists' solution to the problem of purpose has been exactly
>the one suggested by Powers -- selection by consequences. That behavior
>and consequences constitute a feedback system is taken as a basic
>premise....

When Baum says "purpose" he means only "consequence." He doesn't mean that the organism selects a consequence in advance and then brings it about. He is specifically avoiding saying that, because that would put purpose inside the organism and restore the concept that behaviorists have labored so hard to eliminate.

When he says that feedback is taken as a basic premise, he means that the qualitative existence of a closed loop is recognized. He does not use any mathematical or quantitative means of deducing how such a closed loop would work. He speaks of reinforcement as something done to the organism to change its behavior. From nearly a year of correspondence with him, I an certain that he does not want the answer to come out any differently.

Skinner:
>The unit of a predictive science [of behavior] is... not a response but
>a class of responses. The word 'operant' will be used to describe this
>class. The term emphasizes the fact that the behavior OPERATES upon the
>environment to generate consequences. The consequences define the
>properties with respect to which responses are called similar.... The
>OPERANT is defined by the property upon which reinforcement is
>contingent. ...

I'm not sure how this relates to the CT model. It's good that Skinner recognized that consistent ends are brought about by variable means. He clearly defines reinforcement here as a dependent variable, but nowhere else does he treat it as anything but an independent variable.

Why does he have to say that different actions are "similar" when in fact they are different? They may have similar consequences, but that isn't because one action resembles another. It's because of various physical links in the environment that include in their effects a common one. Changing those links wouldn't alter the character of the actions at all, but it would change which actions Skinner would classify as "similar."

The mystery inherent in Skinner's definitions is how an organism can produce actions of many different kinds, even opposite kinds, associated with each other only by the fact that somewhere else, farther along in the causal chain, there is an effect on a common variable. How does the organism know they are going to have a common effect?

> "The only way to tell whether or not a given event is reinforcing to a
>given organism under given conditions is to make a direct test. We
>observe the frequency of a selected response, then make an event
>contingent upon it and observe any change in frequency. If there is a
>change, we classify the event as reinforcing to the organism under the
>existing conditions. ...

I certainly believe in tests. This one, however, doesn't seem consistent with the definition of a reinforcer. A reinforcer is supposed to increase the probability of a behavior on which that reinforcer is contingent. If you make an event contingent on a prior behavior, and that behavior then CHANGES, it would seem that the event is NOT a reinforcer, as it REDUCES the probability of repeating the behavior that preceded it.

On the other hand, if the event now dependent on behavior is observed to change toward a specific form, and on repeated tests with various disturbances converges again and again to the same form because of changes in behavior, there is good evidence that it is a controlled variable and that the organism is acting as a control system with respect to it.

> "The connection between reinforcement and satiation must be sought in
>the process of evolution.... A biological explanation of reinforcing
>power is perhaps as far as we can go in saying why an event is
>reinforcing.

This is premature. It is necessary first to find out whether events have any such power. I suggest that events in the environment have no such extra-physical powers.

> "If we observe someone walking down the street, we may report this
>event in the language of physical science. If we then add that 'his
>purpose is to mail a letter,' have we said anything which was not
>included in our first report?

Yes. We have said that, most likely, any disturbances that would interfere with mailing the letter will result in changes of behavior that will counteract those disturbances and result in the letter being mailed anyway.

>We cannot see his 'purpose' before seeing that he mails a letter, unless
>we have observed similar behavior and similar consequences before. Where
>we have done this, we use the term simply to predict that he will mail a
>letter upon this occasion."

You can't see his purpose even after he mails the letter. A purpose is not just a consequence of behavior: behaviors have lots of consequences, most of them unintended. He may have been reaching in the mailbox to extract a letter, and accidentally let go of the one he had in his hand.

Without applying disturbances you can't deduce purposes at all. Skinner (and Baum) have redefined purpose simply as "outcome." That's why Baum didn't see the vast difference between his explanation and mine. This redefinition is essential, because it allows retention of the basic thesis: behavior is externally controlled.

....

Sorry, Pat. It was Greg who brought up this stuff. I have a very hard time even reading Skinnerian "explanations," much less being nice about it. So I guess I deserve whatever I get.

Best to all,

Bill P.

[from Mary Powers]

re: Beer's bug, agendas, anomalies

I second Gary's idea that Greg could pull the Beer thread out and ship it to him. I do not agree with censoring less temperate remarks (from that loose cannon on the West Coast?). Passion has its place, and occasional outbursts on CSGnet are not out of line. We are not writing for publication here.

*   *   *   *   *

It seems to me that Greg and Pat object to what they perceive as excessive devotion to the CT agenda. Greg's agenda is not purely CT (no problem with this), but includes and is satisfied (to some extent) by CT. It also includes (for example) Bateson and Beer, whom he feels that "pure" CSGers ignore to their loss. Pat, as the resident expert on the net, rightly defends Beer as having some valuable things to say. However, attacking CT as being empirically flimsy, as Greg does, seems to me to miss the point. The point is this: on the one hand there is a generative model (CT) whose empirical data does not at the moment extend to bugs but which hypothesizes certain requirements in bug neurophysiology, and on the other hand there is a presumably empirically-derived artificial bug that works differently from the CT model. I say presumably because the empirical facts built into Beer's bug may not be strictly empirical, but rather based on some implicit and unexpressed assumptions that force the supposedly empirical data to take a particular, non-CT form. Now the question is, whose job is it to figure this out? You seem to assume that it is the theoretician. I think the author of the bug is best equipped to take a further look at his creation while entertaining the notion that it's control systems all the way down, that is, while using an explicit model. Best would be a dialog between the two, intitiated as Gary suggests. If Beer doesn't want to play, no one in CSG will be surprised, but if he does, it could be wonderful.

*   *   *   *   *

There is an article in Science for Feb 7: "When do anomalies begin?" The abstract says:

An anomaly in science is an observed fact that is difficult to explain in terms of the existing conceptual framework. Anomalies often point to the inadequacy of the current theory and herald a new one. It is argued here that certain scientific anomalies are recognized as anomalies only after they are given compelling explanations within a new conceptual framework. Before this recognition, the peculiar facts are taken as given or are ignored in the old framework. Such a "retrorecognition" phenomenon reveals not only a significant feature of the process of scientific discovery but also an important aspect of human psychology.

Further on the authors (Lightman and Gingerich) say:

1) A fact of nature is observed in the context of an existing explanatory framework.
2) The fact does not have a logical explanation in the existing framework but is nevertheless unquestioned and ignored, or accepted as a given property of the world, or simply postulated to be true.
3) A new theory or model is advanced in which the observed fact now has a compelling and reasoned explanation. At the same time, the fact is retroactively recognized as an anomaly in the context of the old theory or model.

The "important aspect of human psychology" mentioned above is cognitive dissonance, and Festinger is quoted: "...when dissonance is present, in addition to reducing it, the person will actively avoid situations and information which would likely increase the dissonance."

The relevance of all this for CT is obvious. The primary anomaly, to me, is the flat out impossibilty of controlled output to work in an environment with disturbances - and its partner, the achievement of consistent ends by variable means. Neither of these appear (that I know of) in the lists that show up now and then of Great Problems Psychologists Need To Solve. CSGers have been publishing (or trying to publish) papers on these "anomalies" for years with little effect. What I'm thinking is that we are tackling the problems that are "taken as givens or ignored", and that the appearance of this article in Science (which is a gloss on Kuhn and may end up being equally trendy), may be helpful if explicitely referred to in future papers. Also, what are the currently fashionable Big Problems in psychology these days, and are they topics CSGers could be addressing productively as a way of getting the camel's nose into the tent?

Mary

Date: Wed Feb 12, 1992 4:05 pm PST
Subject: Re: Behavior

[Martin Taylor 920212 17:30]
(Rick Marken 920211 21:37)
A little time for a more considered response to Rick (and Tom, though I don't expect to quote him).

On re-reading Rick's reponse to my behaviourism posting, I have the feeling

that he and Tom both interpreted my description of what a behaviorist
(I can spell Amurricun, too) approach might achieve as being a description
of what behavioral psychologists have historically done. I get this
impression from:

>You seem to claim that behaviorists (and all other psychologists) determine
>what an organism is controlling (DOING) as a matter of course (because they
>study "behavior" and this implies that they know that behavior is control).

and

>>Behaviourists (depending on the flavour) look the other way around the
>>feedback loop and say that people do what they do because they want
>>something (i.e. they say "the subject picked up thewater glass and took a
>>drink," perhaps adding "to assuage his thirst").
>
>Again, irrelevant. It assumes that we know what "people do". If what people
>do is so obvious, then one should never be able to look at a person's behavior
>and ask "what the hell are they doing?" -- after all, one can SEE what the
>person is doing, can't they? But people ask this about behavior all the time.
>How could this be??? That is what the Mindreading demo is about. You
>know that the subject is moving all 5 numbers -- no question about it -- yet
>observers of the subject's behavior are often inclined to ask "what is that
>person DOING?" ; the program answers the question automatically.
>
My point in the >> part was that the behavioural approach would be to ask
what people DO, in just the same sense as Rick proposes. They would
arrive at the relatively stable construct "getting a drink" based on
"being thirsty", which is not all that far from "controlling the perception
of thirst to a sufficiently low level."

I have always been a bit uneasy about the MindReading demo, in that the
mind-reading computer has access to information not available to the
mystified onlooker--the disturbance. Give an onlooker a control with
which the disturbance could be inserted, and that onlooker might well
"see" which element the subject is controlling.

Before you can apply The Test, you have to have a theory about what options
exist for being controlled, and a way of applying a disturbance that affects
these options (or a subset of them) and not something else that you don't
know about but which might be the thing really being controlled (if
"really" is an applicable word here). Now there are real STATISTICAL
problems (and, Rick, you can't dismiss signal detection theory as you
did, saying it is WRONG--it applies here). Only in exceptional circumstances
can one control a particular percept without at the same time modifying
others that are correlated. This implies that if you do find your
negative correlation with the disturbance, you don't know whether that
was because you had hit on environmental manipulations that transform
into the controlled percept, or because you had got hold of something
that correlated in the subject's perceptual scheme with whatever was
really being controlled. That's problem one.

Problem two arises because it is very rare that anyone controls for only
a single percept at any one time, and conflicts do occur. Furthermore,
control systems have transient characteristics, and complex non-linear
ones have transient characteristics that can depend in non-obvious ways
on their present and historical states. So, if control is not perfect,
and the disturbance is not correctly resisted, you still may have found
the controlled percept. But you wouldn't know you had. You have a problem

both ways.

As an experimenter, you cannot observe what percepts the subject has, so you can't identify which ones are controlled. You can observe only patterns of action that involve things you (experimenter) can perceive. These patterns, in some cases, make sense to you; some aspects of them are repeatable--the car stays more or less in the centre of its lane although you (experimenter) know that there is a variable crosswind. So you try to apply The Test, and do something that would alter the pattern of action if that pattern had something to do with a controlled percept. But it is your (experimenter's) assessment that there exists a pattern, and you test whether the subject also sees and controls something correlated with what you see as a pattern. If you had not seen what you thought might be a pattern, you could not have applied The Test.

Why do you see Testable patterns? Because the environment is not completely unpredictable, and actions that do control certain kinds of percepts do tend to be somewhat predictable. To assuage thirst (or, to lower the thirst percept to a reference level) I do not sometimes pick up a carrot, sometimes bash my head into the wall (that action often appeals to me when I read CSG-L), sometimes put a glass in a puddle, sometimes apply my face to a jet from a tap... I usually find a way to get water into a container from which I can transfer it into my mouth, or I find a mechanism that propels water into a region of space where I can arrange for my mouth to intercept its passage. There is stability in action, labelled in everyday speech by "I'm drinking because I'm thirsty."

There are indeed many actions in the external world that can be used to bring a percept to a reference level. But at the place that you, external observer, can see, there are equally many (actually, orders of magnitude more) variations in the sensor input to the person doing the controlling. Whence cometh the faith that these sensor degrees of freedom are more readily transformed into stable references than are the muscular actions that derive from them? I do not see that either sensors or muscles are determinably closer to the ECS (more probably, banks of ECSs) that accept a transform of one and output a signal that eventually results in the alterations of the tensions in the other.

There are two (at least) central problems: What is observable that could relate to events going on inside the subject, and What class of models describe the observables with least use of special-case variables.

For the second question, I work entirely with the presumption that behaviour is the control of perception. For the first, I am less clear on an answer, and to a large extent what you claim to be observable will depend on the details, or at least the outlines, of the models you use to describe them.

The bottom line is whether there is greater statistical ability to discern patterns by working forward through the sensors to the percept-reference difference, or backward through the muscles to the error-reference difference. There are on the order of 100 degrees of freedom for the joints in the body (by my count), or 600+ muscles (according to somebody posting to CSG). There are millions of sensors, each, in principle, capable of responding differently from its neighbours. Maybe this huge discrepancy in the degrees of freedom at the environmental interface makes it easier to go through the sensors, and maybe it makes it easier to go through the muscles. I suspect it depends on the degrees of freedom in the task (at any moment--the ratios don't change much when you incorporate bandwidths of action or of sensor response). A behaviorist would go through the muscles to look for the stable abstractions that mirror the stable patterns identified in the environment; you would go the other way.

> If what people
>do is so obvious, then one should never be able to look at a person's behavior
>and ask "what the hell are they doing?" -- after all, one can SEE what the
>person is doing, can't they?

No one can't, but one can model it.

Martin

Date: Wed Feb 12, 1992 4:29 pm PST
Subject: To Rick: The "Aw Shit!" Experience

From Greg Williams (920212-2)

My "Holy Shit!" experience regarding PCT came long ago. Since I've been on the net, I've had a number of "Aw Shit!" experiences: when I ask difficult substantive questions of other PCTers and they are ignored. So this time, in hopes of avoiding an additional such experience, I'll make it easy for Rick. One point only, which I should be hard to ignore:

What Beer has done which is more significant than your own modeling efforts is to base his work on experimental data concerning the CONTROL MECHANISM. There are indeed many ways to control in addition to the Powers mechanism of higher level errors altering lower-level reference signals. You seem to think that I mean by this only differences in instantiation (analog, digital, neuronal) -- but I mean ORGANIZATIONAL differences. Some of the circuitry in Beer's bug is constrained by his wanting to mimic real cockroach neural circuit organization. Your models have only a very tenuous relation to experimental evidence regarding real nervous system circuitry. They might match the behavioral data in spades, but that alone doesn't cut it. Beer has a considerably more stringent set of constraints on his models than you. He wants to match neurophysiological data, not just abstract principles.

Greg

Date: Wed Feb 12, 1992 4:31 pm PST
Subject: Re: behaviorism

[Martin Taylor 920212 19:00]
(Bill Powers 920212 12:00)
>
> There's got
>to be some difference between the behaviorist approach and the CT
>approach, or the behaviorist approach would have uncovered the phenomena
>of control long ago.
>
I see them as different kinds of thing. A behavioral approach to CT doesn't strike me as nonsensical, any more than a perceptual approach to S-R modelling. CT depends on the insight that the critical fact in behaviour

is the bringing of some percept to a reference condition.

There are two perfectly adequate and untestable reasons why "the behaviourist approach" did not uncover the phenomena of control (if you exclude JGTaylor) long ago. One is that it takes a fortuitous combination of a prepared person studying the right problem in the right way to generate the insight, and most behavioral psychologists were not trained as engineers (some were), so were not the right people to have the insight. There's some luck there. What would have happened if you had slipped on a banana and decided to study tribology instead?

Second reason may be that you are correct in asserting that it is statistically harder to work backwards around the loop. I don't think that's it, though, because the insight is essentially a philosphical one based on prior knowledge of engineering. As I said some months ago (and you took umbrage), it is trivially true that organisms act so as to control perceptions so that they match references. The insight is that this matters, and provides a hook on which everything else we can study in biology (not only psychology) can be hung. And, to repeat myself from a day or two later, a thing being trivially true does not mean it is trivial to discover or to notice.

I'm not convinced by your demonstration that disturbances enter at several stages between the outermost controlled percept and the patterned event(s) that correspond to the innermost controlled percept.

```
>  NERVOUS SYSTEM -> MUSCLE TENSION->JOINT ANGLE->APPLIED FORCE->POSITION
>                  /                 /            /              /
>                 D                 D            D              D
>
```
>The regularities we observe as behavior occur at the END of this chain,
>not the beginning. If you follow the chain even further outward, to
>"behaviors" such as steering a car, mixing a cake, walking in a figure
>eight, or ice-skating, even more disturbances come into the picture and
>tracing the chain backward becomes even more obviously impossible. In
>every example, the greatest regularity is found at the end of the chain.
>

This is a short piece of the total loop, even if you go as far as the "grand patterned structure" that corresponds to the controlled percept. Maximum stability is in that structure, so far as any other observer can tell. The percept is hidden away inside, and going from the GPS (hey, that IS a well localized system) through the many stages of perceptual transform that also incorporate information from many other PERCEPTUAL disturbances is likely to be pretty dicey as well. I see the kind of focussing you get with a pair of parabolic mirrors facing one another at a considerable distance. The controlled percept is at one focus, the GPS out in the world is at the other. In between, whether the rays are going from the percept to the Grand Patterned Structure or the reverse, they are pretty well dispersed. In the psychological system, there are multitudes of such mirror pairs, the rays from all of them sharing the same intevening space.

Maybe the analogy is still not transparent. But I may try more. It isn't a big deal for me; I'm really only trying to do what Pat was suggesting-- build bridges.

I have a reference level set pretty high for peace and goodwill, but sometimes the way to perceive it isn't clear.

Martin

Date: Wed Feb 12, 1992 6:03 pm PST
Subject: Re: Mary's comments on Beer

From Greg Williams (920212-3)

>Mary Powers

>However, attacking CT as being empirically flimsy, as Greg does, seems to me
>to miss the point. The point is this: on the one hand there is a generative
>model (CT) whose empirical data does not at the moment extend to bugs but
>which hypothesizes certain requirements in bug neurophysiology, and on the
>other hand there is a presumably empirically-derived artificial bug that works
>differently from the CT model.

The foundations (negative-feedback control in the face of disturbances) of PCT are not generative. They do not specify a detailed mechanism for behavioral control. Bill's particular (and highly conjectural) model for how such control might work is generative. As I see it, Beer's bug model satisfies the foundations of PCT, by and large, and therein lies the possibility of bridge-building between Beer and PCTers. The bug model has a GENERATIVE mechanism quite different from that proposed by Bill, and this fact leads to extensive critiques by some PCTers of Beer's modeling. If the clinging to the specifics of Bill's generative model could be let go of (at least to an extent), there would be, I think, recognition by PCTers of the abundant common ground between Beer and PCT.

>Best would be a dialog between the two, intitiated as Gary suggests. If Beer
>doesn't want to play, no one in CSG will be surprised, but if he does, it
>could be wonderful.

Initiated as Gary suggests, with the addition of your suggestion that "loose cannon" remarks not be deleted, I certainly wouldn't be surprised if Beer doesn't want to play the game. After all, how much input to the rules of the game would he feel he has? Nevertheless, it's worth a try, in my opinion -- for the PCT agenda! As I posted privately to Gary, I think Gary is the one best situated to pull out stuff from CSGnet and send it to Dr. Beer. I'm fully occupied with typesetting LCS2 right now, and, of course, am emotionally involved in the debate.

Greg

Date: Wed Feb 12, 1992 7:10 pm PST
Subject: chestnuts roasting...

[from Joel Judd]

************WARNING: THIS IS A TESTIMONIAL**************

Ah, it's good to know when it's snowing outside in Illinois, I can curl up in front of a terminal and bask in the heat generated from discussions on the net.

For what it's worth, I bring considerably less sophistication to any discussion of living systems than any of the regular participants on the net--virtually none at all to discussions of neurophsiology of cockroaches (personally, my wife and I are tired of the adaptable little buggers in our kitchen). But in spite of my naivete, I would just like to say why I persist in antagonizing people in SLA with this control theory.

I was dissatisfied with explanations (and my understanding) of teaching and learning before Gary suggested I look at PCT. I have been teaching for 7 years now and am STILL not happy with what I think I know about teaching and learning. Certainly PCT has little to directly apply at present. Yet I am going to gradute (I think) with a dissertation spouting all kinds of strange things about control of linguistic perceptions. PCT appeals to me not because of what it lays out in black and white but because of what it IMPLIES for future understanding. It is not too specific, perhaps, in terms of models or empirical demonstrations, but it is more specific than recent ruminations by other more widely read theorists and researchers in psychology. And that's the second reason I find it satisfying. For example, in spite of what Baum says (actually BECAUSE of what people like Baum say), Jerome Bruner spends the better part of a book (_Acts of Meaning_) showing that most psychologists/cognitive scientists still eschew INTENTION or PURPOSE in psychology. And Mark Bickhard gets perilously close to a hierarchy of perception, but seems to lack the notion of a closed loop which would make his model work. So I find PCT giving me a greater understanding of works I may have read five years ago, but didn't understand; or, works which gave me a funny feeling that something was wrong which I couldn't put my finger on.

In short, PCT may leave a lot to be desired, but it's a helluva lot better than the alternatives *I* have had to work with so far (how's that for falsificationist sentiment?).

Date: Wed Feb 12, 1992 9:47 pm PST
Subject: Re: Beer and Behavior

[From Oded Maler (920212]

concerning [Rick Marken (920210)]

The first part of your post was a very nice explanation of the difference between possible ways to describe behaviors. However, the rest is a bit weak, in the sense that it repeats things that I assume are obvious to anyone who tries to think about robotics seriously (at least it is obvious for me). For example:

>so if you want the bug to control a particular variable you know that you
>need to design a way for the bug to perceive that variable; you also know
>that you need a way for the bug to affect the variable in all dimensions
>in which the variable is to be controlled.

For simple sensor-effector loop this is an obvious "motherhood" statement. For complex loops this is still true but not very prescriptive.

later you say:

> you have to ask "what variables
> does the robot need to control for ITSELF in order to produce a particular

> observable behavior for me"?

Recall that after all, inside the robot there is nothing but circuitry relating sensors, effectors and intermediate signal. You are free to interpret some of them as references, errors etc., as you have done with Beer's model. The point is that for a simple system, such as obstacle avoidance, you can do it either using the PCT terminology or without it (as most roboticist, real or simulated, can do). This is not a big deal. The questions is whether for a much more complex kind of behavior, with many sensors and effectors, with complex interaction with the environment, the PCT 'methodology' will help to find solutions much more easily than without using it. A-priori it seems that the problems of engineering creativity are delegated within this framework into the problem of defining the perceptual variables and the compensation methods for them. Whether or not this makes life easier, I am not yet in a position to judge.

> Moreover, Bill's "Little Man" demo
>shows (in beautiful detail) some of the complex, coordinated behavior
>that the bug shows -- and again, the basic principle of design is
>quite clear.

So let me ask for that demo.

>Again, if the goal is to produce amusing "behavior" -- ie,
>results of action, then perhaps PCT could get some notice by building
>ping-pong playing robots; I bet it would not be too hard to give the
>"Little Man" another arm and have him juggle instead of just point.

If you succeed easily in making such a thing in a realistic simulation, including noise in motors and vision systems, uncertainty in results of collisions, and real-time constraints - I'll bet you could publish it rather easily. I'm also sure that if you make one of Brooks' *real* artificial insects walk smoothly in a rugged unstructured terrain, using PCT principles, PCT will become a very hot and respectable topic (which is of course not what you want.. [p.s. it always reminds me about the difference between Judaism and Christianity in their approach to marketting Monotheism, so I somehow have a basic sympathy to your approach, but let's not get into this can of worms..])

Just remember that inside a robot, at least, there are just circuits and interrelated signals. You will not get far if you go to someone and say "you should interpret the circuit this way and not the other", or "the way you use the word 'behavior' is incorrect" - this is silly if you read some of the lengthy discussions about words and meaning.

Best regards

--Oded

Date: Thu Feb 13, 1992 6:31 am PST
Subject: FTP - correction

A minor correction:

>If you have not used ftp before, here is a guide to getting this file:
>
>cd csg

```
>binary
>get dem2a.exe  (This copies the file to your machine)
>quit
```

The command is "cd pub/csg", not "cd csg". However, if you use the server, you can omit the "pub" directory, since it knows all about it.

To obtain more information about the mail server send the message "help" to server@biome.bio.ns.ca. If you want information on how to use ftp, send the message "ftp" to the server. You can send both commands in one message, but on separate lines.

Date: Thu Feb 13, 1992 6:32 am PST
Subject: COCKROACHES AND THEIR BEHAVIOR

FROM CHUCK TUCKER 920211

I am the cockroach killer in my household and over the years I have made tests to see what they were controlling for so I could accomplish my task of killing them with greated efficiency (see Greg was correct; PCT can be used in a harmful manner) and although my test are not very systematic I have found that my roaches control for light. They appear to be controlling for and edge because it is dark; just shine a light on them and watch them run. Find a roach in the middle of the room (this is not easy to do), turn the lights on, put a shadow over them and they will stop; then step on them. On PBS on Tuesday they had a program on smells and they showed how roaches controlled for the scent of virgin females. They also demonstrated that roaches make tracks to apparently mark their trail like other animals do.

I have not read any of Beer's work or seen the BUG program; these are just the observations of a long time roach killer.

Regards, Chuck

Date: Thu Feb 13, 1992 7:32 am PST
Subject: dis(covering)agreements

[From: Bruce Nevin (Thu 920113 07:32:57)]

(Martin Taylor 920212 14:30) --

I think (and you have suggested too) that we have more agreements than we sometimes let on. :-)

I agree that all three points of view (social history, individual learning/acquisition, access for individual use in social settings) address real contributors to what and how it is (where "it" is language, social convention, etc., pick one). I do not urge the importance of one exclusive of the others. I do not yet have a good sense of the relative importance or weight of their contributions, or even of how to guage that.

I suggested the primacy of spoken language over written records for understanding language in PCT terms. When I do this, I am speaking from the point of view of the individual learning and using language in the present, not the ancient history of writing systems.

The point is relevant to CSG insofar as examples people have offered have usually omitted aspects of spoken language that are not recorded in English orthography. It is easy to imagine that word recognition is a task of recognizing event-sequences exactly analogous to recognizing strings of letters. With that image in imagination, mention of syllable structure, intonation, and so on, may seem abstruse and academic. I am open to the possibility that perception of segments is a byproduct of higher-level analysis, as has been argued for perception of individual windows in a wall of windows or individual Ts in a string of Ts, an argument that the configuration containing them is perceived first. In this case, the syllable or semisyllable may be the unit of perception, only analytically broken down into constituent segments. (Has anyone tried to teach reading by teaching the alphabetic representations of syllables or semisyllables?) But I agree for that relevance extended discussion is unwarranted, only occasional reminders.

I agree that the interaction between spoken and written forms is subtle and complex, and said as much. I only argue that spoken language comes first--not only temporally first in terms of normal language acquisition, which is indisputable, but also primary in the importance and pervasiveness of its influence in the perceptual control of language. This is less obvious, and you may choose to debate it off line. But please understand that I do not thereby deny visual iconicity, ranging from the breast-like golden-arch M of MacDonald's to e. e. cummings's exquisite mini-haiku lyric

l(a
le
af
fa
ll
s)
one
l
iness

It is not quite possible intelligibly to read this aloud, even if you give the phrase that interrupts the word a much reduced intonation contour. But once you "get" it, the ambiguity of letter l with numeral 1 (especially on cummings's typewriter), the "one" and the "I-ness," and even the falling-leaf-like swirling of the letters down the page, all help to articulate the associative meanings of the spoken words, and even the just-out- of-reach inability to read it satisfactorily aloud to someone helps to evoke the intended emotional ambiance. Rebus-like, don't you think?

(Martin Taylor 920212 14:40) --

It seems to me that there is no need to introduce disturbances into this ATC system, but rather to determine how to instrument the disturbances that are already present. This seems more realistic psychology, if the disturbances are not under the experimentor's control but rather must be deduced, teased out of the situation at hand, together with the disturbed variables (perceptions) that the active participant is controlling.

More on this re empathetic use of imagination, below.

(Mary Powers) --

> a way of getting the camel's nose into the tent

This is why it helps to approach adversaries by way of points of agreement rather than points of disagreement. Not to convert them, but to provide a situation in which they can be exposed to control theory. And the point is not to pretend agreement (that's what the dismissive Behaviorist reviewer does, saying that every CT X reduces to a Y of S-R theory) but to re-frame the point of agreement with respect to an internal contradiction that the S-R advocate has overlooked because it occasions cognitive dissonance. And what is that cognitive dissonance? What controlled perceptions are being disturbed (unless he or she ignores something)? What control of the same or alternative perceptions can be immediately offered to assuage the discomfort?

Brava, Pat, for bringing this up again, and so cogently!

(Martin Taylor 920212 17:30) --

>I have always been a bit uneasy about the MindReading demo, in that the
>mind-reading computer has access to information not available to the
>mystified onlooker--the disturbance. Give an onlooker a control with
>which the disturbance could be inserted, and that onlooker might well
>"see" which element the subject is controlling.

Wouldn't you want the onlooker to have a display of the correlation between action and disturbance for each number, with an understanding that a "fairly strong negative correlation between action and disturbance" was an indicator of control? This is the information that the program has. Providing such a display as a final phase of the demo, Rick, would underscore the point that without that information you literally don't know what you are talking about when you use words like "behavior." Just looking at behavioral outputs, you don't know what's significant. An everyday experience, in fact, but one that we get ourselves out of as quickly as possible, if necessary by strong use of imagination and our capacity to ignore and rationalize--too much cognitive dissonance!

>Why do you see Testable patterns?

I think the usual answer is empathetic use of imagination. I intuit what perceptions I would be controlling were I in the shoes of the participant that I am observing. Intuition and empathy of course need verifying. We all know two methods preferred here are (1) to disturb the supposedly controlled perception and see if the participant does in fact counteract the disturbance, and (2) to model the participant-situation-disturbance system and compare with observations by method (1). But I think the ideas come from empathetic identification with the observed control system, in imagination. Maybe we augment that with self-observation under analogous circumstances. This all rests on an ability to create and use analogies. These in turn hinge on use of categories and often (always?) on use of language.

(Greg Williams (920212-2)) --

>There

>are indeed many ways to control in addition to the Powers mechanism of higher
>level errors altering lower-level reference signals.

Bill has made the point that the phenomenon of control is the central matter, that the levels of the hierarchy or the architecture of the ECS or the existence of a hierarchy of ECS-like mechanisms, all of these details of implementation are expendable, if a better way of modelling control comes along.

I have argued against social "control" on grounds that there is no known mechanism analogous to the transmission of neural signals between comparators by way of I/O functions that combine them. This argument is weak just because alternative means and architectures for implementation are possible. I have mentioned pheromones, pervasive neuropeptides, and empathetic use of imagination to learn social norms and conventions (exactly analogous to the experimentor's use, above).

A still weaker argument that I have suggested is that the reference signal is obligatory for an ECS, not for a person in a social setting. This is weak because individual control for conformity to social norms is unconscious to the extent that it is authentic, not a matter for conscious choice whether to conform or not, and therefore to all effects just as obligatory. We become aware of social conventions only in proportion to our not participating fully in them. And disturbance to the perception of someone's conformity to a convention may not result in counteraction to resist the disturbance, but rather a shift in category perception as to what kind of person it is. The same must happen sometimes with intrasomal perceptions that are diagnostic rather than being controlled strictu sensu.

Martin, it is true that I have backed off from these issues since I first joined the list last March or April, but only because other concerns had higher priority for me. And if you perceived me as castigating you about this, I apologize. I recall responding to you with the above two counterarguments re social "control."

Bruce
bn@bbn.com

Date: Thu Feb 13, 1992 9:03 am PST
Subject: Building a better bug

From Pat & Greg Williams (920213)

>Oded Maler (920212)

>If you succeed easily in making such a thing in a realistic
>simulation, including noise in motors and vision systems, uncertainty
>in results of collisions, and real-time constraints - I'll bet you
>could publish it rather easily. I'm also sure that if you make one of
>Brooks' *real* artificial insects walk smoothly in a rugged unstructured
>terrain, using PCT principles, PCT will become a very hot and
>respectable topic...

We don't think PCTers have to do THAT much. Beer's bug is "externally" robust, but NOT very robust (in some places) "internally." E.g., the parameters of the bug's neurons and their connections are, in some cases, quite "touchy" --

change them a little and control disintegrates. Quoting from Beer's book, p. 167: "Even the biologically-inspired controllers had to be fine-tuned by trial and error (there are over 500 parameters in the artificial insect's nervous system!)." In private communication with us some time back, Beer implied that the fine-tuning was a real pain -- one of the reasons he wanted to investigate organization of circuits with GAs. If PCTers could make a bug model (still using neuron models like Beer's and incorporating a modicum of empirical findings) with significantly enhanced internal robustness, we think that the robotics world, at least, would be enthusiastic.

Pat & Greg

Date: Thu Feb 13, 1992 9:40 am PST
Subject: Re: dis(covering)agreements

[Martin Taylor 920213 12:15]
(Bruce Nevin 920213 0733)
>
>
>I think (and you have suggested too) that we have more agreements than
>we sometimes let on. :-)
>

Must be so. I really couldn't find anything in your posting today to disagree with. I must be getting senile or something :-). I liked everything you said about written and spoken language, about using empathetic imagination to guess what might be controlled perceptions, about Rick providing a correlation display ...

What's going on here?

Martin

Date: Thu Feb 13, 1992 11:55 am PST
Subject: Beer and PCT

[from Gary Cziko 920213.0845]

Greg Williams (920212-3) says:

>I think Gary is the one
>best situated to pull out stuff from CSGnet and send it to Dr. Beer. I'm fully
>occupied with typesetting LCS2 right now, and, of course, am emotionally
>involved in the debate.

It wouldn't be hard for me to pull out the Beer-related posts and send them to him and I may do this when (if?) the topic winds down.

But even though it was I who first suggested this, I don't think that it will have much effect since I believe Beer sees himself already doing the types of things PCT would suggest plus other stuff as well.

I also want to add that from my conversation with Beer I believe he has both biological and engineering interests (indeed, he has appointements in both biology and computer science). One stream (the biological) of his work wants to use everything we know about the nervous system and behavior of bugs and try to recreate both via computer. The other (the engineering) stream involves using whatever ways one can to get the bug to do what you want it to do using such techniques such as genetic algorithms to design leg and gait "controllers." For the latter purpose, you may end of with circuits that work but nobody can figure out how or why they work (cognitively inpenetrable). But why would an engineer worry about that? (A leading question if I ever wrote one).--Gary

Date: Thu Feb 13, 1992 12:25 pm PST
Subject: MindReading Program

[from Gary Cziko 920213.1345]

(Martin Taylor 920212 17:30) --

>I have always been a bit uneasy about the MindReading demo, in that the
>mind-reading computer has access to information not available to the
>mystified onlooker--the disturbance. Give an onlooker a control with
>which the disturbance could be inserted, and that onlooker might well
>"see" which element the subject is controlling.

The MindReading program does not need to have access to the disturbance to determine which numeral is being controlled by the subject. Instead of calculating the correlation between disturbance and the mouse movements, and picking the numeral whose disturbances show a high negative correlation with the mouse movements, it could also calculate correlations between the mouse movements and the movement of each numeral on the screen. The numeral with the LOWEST correlation with mouse movements is the numeral being controlled.

Rick changed the program to do just this at the last CSG meeting in Durango for me after being recommended by Bill Powers.

What I find so intriguing about the MindReading program is that the computer can find out which number your are controlling more or less regardless of what type of movement pattern you are controlling for. Whether you keep one number still, move it in a circle, or write your name, the program can find the controlled numeral. In fact, all you need to do is have a plan of where you want the numeral to go next. Because if you have a plan (purpose, intention), you will resist disturbances and the correlation between the controlled numeral and your behavior will be lower than for the uncontrolled numerals floating on the screen.--Gary

P.S. Rick, you should think of spiffing up these programs (including, for example, variable speed for use with Macs of varying speed) so we can make them available via the file server.

Gary A. Cziko Telephone: (217) 333-4382

Date: Thu Feb 13, 1992 2:29 pm PST
Subject: Stability of remote effects

[From Bill Powers (920213.1300)] (Resent w/ revised format 7:12 pm)

Martin Taylor (920212.1730) --

RE: getting thirsty

This concept doesn't fit my stereotype of behaviorism. I had thought that behaviorists eschewed references to unobservable inner states. In place of "thirst," hours of water deprivation would be substituted. Even a "stimulus" or a "percept" would have to be defined in terms of something observable from outside the organism -- that is, something is a stimulus object only if there is a response to its presence. Has this sort of requirement been abandoned?

RE: Mind reading

>I have always been a bit uneasy about the MindReading demo, in that the
>mind-reading computer has access to information not available to the
>mystified onlooker--the disturbance.

You are right. It is possible, however, to do the mindreading experiment without access to the disturbance. The principle then involved is that among all the variables known to be affected by the actions, the one showing the _lowest_ correlation with the action is (most likely to be) the controlled variable. This requires only knowing the nature of the connections between the action and each potential controlled variable, which in RIck's demo is simple and the same for all objects (numerals floating about on the screen). This method isn't used in the demo because the variances are larger and it takes much longer to get a reliable measure of the lowest correlation. If the subject picks a very simple pattern of motion for the controlled numeral, looking for the lowest correlation works better than with complex patterns -- best of all if the participant just keeps the numeral in one place. The simpler the pattern, the easier it is for the uninformed onlooker as well as the computer to figure out what's being controlled without knowledge of the disturbance.

>Before you can apply The Test, you have to have a theory about what
>options exist for being controlled, and a way of applying a disturbance
>that affects these options (or a subset of them) and not something else
>that you don't know about but which might be the thing really being
>controlled (if "really" is an applicable word here). Now there are real
>STATISTICAL problems ...

The general application of the test, in which the reference level for the controlled variable isn't reasonably constant, does get into statistical problems, and all your remarks apply. The human observer, however, can do much that statistical tests can't do. If there is any pattern in the variations of reference signals, they will be reflected in the observed behavior of the controlled variable. The observer, being human and being familiar with many of the patterns people are likely to produce, can often see a regularity which can then be incorporated in the definition of the controlled variable (for example, if there is a suggestion that one or more of the floating numerals traces out a square). When that is done, the statistical measures would become very much clearer.

It's well to keep in mind that real applications of the Test don't rely only on applying random disturbances, and they aren't confined to looking for statistical maxima and minima. Systematic disturbances based on the hypothesized nature of the controlled variable can discrimate much more precisely than can random ones. If, after several candidates have been found through looking for minimum or maximum correlations, there are still uncertainties of the kinds you mention, the next step is to try to block perception of each candidate. The one or ones that are actually under control, and only those, will cease to show signs of control when it or they aren't perceived. And finally, the connection from the action to each potential controlled quantity can be investigated -- if the action isn't actually connected to the controlled variable (if something else is making the variable change), the participant can't be controlling that variable or any function of it.

The ultimate test is (when possible) to make a model based on the hypothesis and compare its behavior with that of the subject, quantitatively. If the correlations between model behavior and that of the participant are in the neighborhood of 0.95 and better, and are far lower for the other candidates, you can be pretty confident that you have identified a control system and its controlled variable.
The Test is really just a strategy for finding variables that are under control: it isn't a cookbook method. Statistical measures can provide you with a lead and narrow the field of possibilities. But given that lead, you can then try systematic approaches that narrow the possibilities further. You may fail. But when you succeed, given high standards for accepting the appearance of success, there will be no reasonable doubt that you understand what is doing on.

One of my objections to the statistical approach to understanding behavior is that after the first significant statistical measure is found, the experimenter quits the investigation and publishes. If you get a correlation of 0.8, $p < 0.05$, you next question should be, "Where is all that variance coming from?" If you set your sights on 0.95, $p < 0.0000001$, you won't quit after the preliminary study, but will refine the hypothesis until you get real data.

I think it would be a good idea if you were to do some control-system experiments, even very simple ones. It's hard to get a feel for the way the control model matches behavior without going through the details for yourself. I think many of your cavils and cautions are based on a rather abstract picture of control relationships, which would be much sharpened by participating in control experiments and fitting some models to the results. I know that stick-wiggling is not a very interesting kind of behavior, but until you've experienced it and seen the results first- hand, you can't appreciate how different the results are from those obtained in just about any kind of "normal" behavioral experiment you could name.

Are you aware of, and have you tried, the Coin Game? This game illustrates many of the features of the Test, as well as many of the problems (which you have correctly foreseen). The main thing it reveals is the step-reduction in uncertainty that occurs when, as experimenter, you finally hit on the correct definition, or one very closely related to it. Not only do you understand what is being controlled, but you understand every move theother person made while you were trying to find out what is controlled.You also discover how you can be misled by false-positive results that seem, for a while, to support your hypotheses, and you come to appreciate the devastating effect of a single clear counterexample.

If you haven't heard of this Coin Game (it's described in BCP), I'll describe it on the net.

>Whence cometh the faith that these sensor degrees of
>freedom are more readily transformed into stable references than are
>the muscular actions that derive from them?

I don't get what you mean by "transformed into stable references." I don't see at all how muscular actions would be transformed into references, stable or not.

RE: stability at end of causal chain

>This is a short piece of the total loop, even if you go as far as the
>"grand patterned structure" that corresponds to the controlled percept.

Suppose that the "grand patterned structure" is something simple like the position of a car on the road. There are indeed countless percepts involved in perception of this position, but at some point in the hierarchy there's a signal standing for that position -- else you wouldn't be perceiving it. As the car drifts from side to side, so does this perceptual signal change state (without sticking my neck out as to how).

The perceived state of the car is compared with the intended state -- with a signal specifying how the perception should look. The error is derived from the difference between perception and reference (I know you know this, I'm just going through the steps).

The error is transformed, eventually, into muscular effort applied to the steering wheel. This effort cocks the front wheels and causes a sideward force on the car.

Also acting, at the same time, are independent external disturbances from crosswind, soft tires, nonlaminar airflow, bumps, and tilts in the roadbed. These, too, apply sideward forces to the car, independently, unpredictably, and invisibly. The actual path of the car results from the sum of the steering effects and all these other disturbances.

The actual path of the car is found to stay within a foot or two of the centerline of a straight road on a trip extending 20 miles: an average directional steering error of at most 4 seconds of arc.

Given the observed behavioral result, we could now try to explain it by reasoning backward or forward. Going forward, we would say that the position of the car was continuously perceived with an accuracy of plus or minus 2 feet and compared with a reference position corresponding to the center of the road (or we could apportion the error between reference signal and perceptual signal). This error drove the steering efforts, with some error sensitivity and some dynamic stabilization. The control model then predicts the observed behavior, given only that the loop gain is high enough to limit the effect of the maximum total expected disturbance to a deviation of two feet. Or we could reason backward, looking for the stimulus that caused the efforts that resulted in the observed position of the car. This stimulus might be the error signal in the same model. So how do we work backward?

First we must know how the car responds to sideward forces. Then we must plug the record of those forces into the dynamical equations to deduce what the total disturbance was at every moment. Knowing that, we can see how much added steering force was required to keep the car within two feet of the center of its lane during the trip. Then, by acceptable hypothesis, we could say that this required force is a measure of the error signal that was needed. We have now arrived backward at the same point we reached going forward through the perception and the comparator.

The claim is that there's no appreciable difference between going forward and going backward to reach the error signal. There is in fact an enormous difference, five or six orders of magnitude.

To go backward, you would require a complete record of every disturbance that acted on the car during the trip. How accurate would the data have to be? I figured it out once, and it came to something like one or two parts in ten to the eighth power. This is because in computing the path of the car from the forces acting on it a double time integral is involved. You would have to know, for every instant, all the forces affecting the car, and know them with this accuracy. The steering efforts would have to be deduced with the same accuracy, so that the muscle tensions and the driving neural signals could be deduced with the same accuracy.

Even worse: you would now have to claim that the variations in error signal just deduced would in fact be transformed through the nervous system, the muscles, and the steering linkage into a force on the car, adding that force to extraneous ones, with an accuracy of a few parts in 10^8. This is clearly far beyond the bounds of reason. The actual achievable accuracy in converting a neural signal to a force is something like 5 percent - 5 parts in 10^2.

In the above, I have assumed that you have access to essentially perfect measures of each source of disturbance throughout the trip, and have a perfect theory that converts these disturbances into equivalent forces on the car. In practical terms, that is impossible, not only quantitatively but qualitatively. And we are speaking only of analyzing the results of one car trip after the fact.

Prediction is out of the question. In order to predict the path of the car on the return trip from knowledge of the error signal, you would have to be able to predict the crosswind (that is, the weather) with an accuracy of a few parts in 10^8 all during the trip. You would have to know the location and size of every bump in the return lane, the exact tilt of the roadbed in that lane, the rate of loss of air from each tire, and so on -- all with an impossible accuracy.

So when I indicate those disturbances in the chain of cause and effect, I'm not just appealing to generalities. The disturbances are real, and theycan't be sensed with enough accuracy by even a billion dollars' worth of instruments to permit their effects on the final outcome to bepredicted well enough to allow backward tracing of causes. The nervous system is totally incapable of that accuracy, and it doesn't even have sensory access to most of the causes of disturbance.

Yet it is the outcome far down this causal chain that is controlled with great precision by the nervous system. This result is clearly, obviously, not accomplished by computing the signal required to achieve the final result. Any command-driven model of behavior is untenable as an explanation of the observed constancy of behavioral effects remote from the nervous system. Reasoning backward from these effects to the neural signals that brought them about is impossible. Creating neural signals by central computation that will have such precise and repeatable results is impossible.

Martin, you probably feel that you turned on a faucet and got hit with a firehose. I apologize for any overkill.

This was simply an opportunity to put forth my argument in full detail. My argument is simple: when you look at the physical details involved in producing any regular behavior at all, it becomes clear that no model in which driving neural signals are precalculated, or produced in response to stimuli without immediate concurrent negative feedback, can come within 5 or 6 orders of magnitude of explaining what we observe. Even without considering control theory, both the S-R model and the cognitive model are refuted by the facts of real behavior in a real world. And when all theories that involve reliance on a regular output chain are eliminated, there is only one theory left that fits the facts: control theory.

Is all of this common knowledge among behaviorists?

Your turn. --

Best,   Bill.

Date: Thu Feb 13, 1992 2:54 pm PST
Subject: Beer bridges, behavior

[From Rick Marken (910213)]

Greg Williams (920212) says:

>   I'll make it easy for Rick.
>One point only, which I should be hard to ignore:
>What Beer has done which is more significant than your own modeling
> efforts

Ouch! I hear ya!

>is to base his work on experimental data concerning the CONTROL
>MECHANISM. There are indeed many ways to control in addition to the
>Powers mechanism of higher level errors altering lower-level >reference
signals.

No doubt.

>   You seem
> to think that I mean by this only differences in instantiation (analog,
> digital, neuronal) -- but I mean ORGANIZATIONAL differences. Some of the
>circuitry in Beer's bug is constrained by his wanting to mimic real cockroach
>neural circuit organization.

Yes. I know. But I think the basic organization of neural circuits must be a negative feedback control loop. There must componenets that sense, and components that effect. There must be a cmoponent like a comparator -- even if it has no explicit reference input. I think something must serve as a comparator because there must be a way of converting sensor signals into effector actions. There can be gain modifiers and all kinds of things influencing the loop -- but ultimately it must be organized as a negative feedback control loop if it controls.

>   Your models have only a very tenuous relation to
>experimental evidence regarding real nervous system circuitry.

Well, it's not THAT tenuous. But I'll concede it.

>   Your models might match
>the behavioral data in spades, but that alone doesn't cut it.

Depends on your goal. But, I basically agree. I would like the model to have some reasonable correspondance to what we know of neurophysiology too.

>Beer has a considerably more stringent set of constraints on his models
>than you. He wants to match neurophysiological data, not just abstract
>principles.

He may be trying to stay within physiological contraints as he sees them. But I don't recall seeing anything in the book or the article that could be called "matching his data to the neurophysiological data". What data does he want to match? Has this data been obtained from a real "working" bug. If not, why build a whole bug? If this data was collected from a "working" bug, was it collected from a bug that lived in the same kind of environment as the simulated bug?

I've glanced through the book again. The closest thing I find to a comparison of simulated to real bug behavior is on p. 136 and 137. He reports a general correspondance between level of satiation in the simulated bug and interbite interval data in hungry Aplysia. He does not explain the graphs well enough for me to know what they are showing. The Aplysia data is average latency data (whatever that is) and the bug data is inter-bite interval for a single bug. He presents these graphs to show that the interbite behavior of the bug during feeding depends on satiation in about the same way as the interbite behavior of Aplesia depends on how long they have been without food. There seems to be a family resemblance between the curves. But there is no quantitative analysis to show how well the model matches the phenomenon.

But this level of match seems to be all that Beer is looking for -- for example, on p. 159 of the Discussion chapter he says "P. computatrix displays a number of interesting behavioral characteristics which are strickingly reminiscent of natural animals". I think the same could be said of the "Gatherings" program. And considerably more could be said for Bill's "Pigeon conditioning" program in which a simulated pigeon produces exactly the exact response rate data that is found in conditioning experiments (under many different schedules). Of course, Bill's model does not match the neurophysiology as exactly as Beer's presumably does. But again, I have looked through Beer's book and I find absolutly no comparison of the physiological simulations in the model with physiological data obtained from actual behaving bugs. Beer's ostensible goals are the same as the PCT goals (as you and Gary and others have pointed out) -- he really just wants to mimic BEHAVIOR. If that is his goal, then I think I would argue that PCT has achieved that goal in many instances and with far greater precision. I don't see this kind of success in Beer's models. Now it may, indeed, be that we have much to learn from Beer's attempts at neurophysiological accuracy. But, quite frankly (and I am not dismissing Beer's work -- this is just my opinion) I don't see what it is. I'm quite willing to accept that this is because I am stupid or ideologically blinded or whatever. But I need help on here. I really don't understand what I'm supposed to be learning from Beer.

But this post is supposed to be about building bridges. So I will admit (as I have many times before) that Beer is building systems that control variables. I hope he will take a look at PCT as a possible way of getting a deeper understanding of what he is doing. As my start at building my side of the bridge toward Beer, I plan to look at his book and try to figure out how his gait generation system works -- so I can see how PCT might look at it.

Greg Williams says to Mary Powers

>As I see it, Beer's bug model satisfies the foundations of PCT, by and
>large,and therein lies the possibility of bridge-building between Beer
>and PCTers.

I concur -- see paragraph above.

>The bug model has a GENERATIVE mechanism quite different from that
>proposed by Bill, and this fact leads to extensive
>critiques by some PCTers of Beer's modeling. If the clinging to the specifics
>of Bill's generative model could be let go of (at least to an extent), there
>would be, I think, recognition by PCTers of the abundant common ground
>between Beer and PCT.

Point taken. However, I didn't object to Beer because of the generative mechanisms in the model. It just seemed, from reading the book, that he doesn't have a good grasp of the nature of behavior. But, as I said, I will look at Beer's gait generation system to see what is going on from a PCT perspective -- if I can understand it. Actually, since you and Pat wrote the program, maybe you could explain it? I would really like a simple explanation. The bug's gait is very interesting -- I wonder if there is a system for generating such a gait with input control rather than output generation (with input control at each individual leg). I'm reaching across the stream (literally, here in LA) -- I hope there are folks on the other side willing to build towards me.

Martin Taylor (920212 17:30) says:

>I have always been a bit uneasy about the MindReading demo, in that the
>mind-reading computer has access to information not available to the
>mystified onlooker--the disturbance.

The point of the demo is that you can look at behavior and not know what a person is doing -- even though you can see what they are doing. As I noted, this happens all the time in "real life". Moreover, the disturbances that are being resisted during "real life" behavior (forces such as those of gravity and angular accelerations created by the actor him/herself, etc) are typically invisible. The mind reading demo is just an attempt to get people to think about what it means to say that someone is "doing" something.

>   Give an onlooker a control with
>which the disturbance could be inserted, and that onlooker might well
>"see" which element the subject is controlling.

In fact, WILL SEE the controlled variable. I have the computer do it just for dramatic effect. You actually suggested a great demo -- though it takes two mice -- the second mouse is used by an onlooker to disturb each number in turn (by pushing the mouse button the disturbance could be switched from one to the other number). I bet the onlooker, without any statistical test, could "feel" the resistence when the disturbance was applied to the controlled number. This is what "the test for the controlled variable" is about. The onlooker is a control theorist (or any one savvy enough to know that people are controlling variable aspects of their experience) doing the test for the controlled variable. In "mindreading" as it exists, the computer is the control theorist, doing the test.

>Before you can apply The Test, you have to have a theory about what options
>exist for being controlled,

I'd call it a hypotheses -- right. In the mindreading game I've sort of
limited the hypotheses to 5 -- though I did do a version where you
could also control the distance between certain pairs of numbers -- so that
was another possible set of hypothesis.

>   and a way of applying a disturbance that affects
>these options (or a subset of them) and not something else that you don't
>know about but which might be the thing really being controlled (if
>"really" is an applicable word here).

Right -- the test is easier said than done.

>   Now there are real STATISTICAL
>problems (and, Rick, you can't dismiss signal detection theory as you
>did, saying it is WRONG--it applies here).

SDT is fine for radar. I didn't say it was wrong. I said the people who
use it have the same, incorrect concept of behavior as any other psych-
ological theorist. Behavior is a dependent variable of their choice. SDT
people don't test to see what people are controlling -- do they??

> Only in exceptional circumstances
>can one control a particular percept without at the same time modifying
>others that are correlated.

My point, exactly.

> This implies that if you do find your
>negative correlation with the disturbance, you don't know whether that
>was because you had hit on environmental manipulations that transform
>into the controlled percept, or because you had got hold of something
>that correlated in the subject's perceptual scheme with whatever was
>really being controlled. That's problem one.

In principle you can eventually nail it down. But if you don't do the test
your guesses about what a person is doing are just that -- guesses.

>Problem two arises because it is very rare that anyone controls for only
>a single percept at any one time, and conflicts do occur.

To paraphrase one sage's deathbed observation about comedy -- "dying is easy,
science is tough"

Oded Maler (920212] says:

>Just remember that inside a robot, at least, there are just circuits
>and interrelated signals. You will not get far if you go to someone
>and say "you should interpret the circuit this way and not the other",
>or "the way you use the word 'behavior' is incorrect" - this is silly
>if you read some of the lengthy discussions about words and meaning.

My point is that the behavior of living systems consists mainly of the
control of perceptual variables relative to internally specified reference
levels. All consistent behaviors are controlled variables -- from leg
positions to political positions. If you want to design a system that
behaves like a living system, that is controls, then you must build it to
be organized around the control of input. Complex controlled inputs, like
brisses and communions, require systems that can perceive these results.

I know that this will not solve your problems as a roboticist but if you understand it, you can start working in the right direction -- ie, trying to build a briss perceiving system rather than a system that produces particular actions that produce a "briss". I want you guys to do the hard work -- but if you are looking at all toward the life sciences for clues about how to build "artificial living systems" I think your best bet is PCT.

Hasta Luego

Rick

Date: Thu Feb 13, 1992 4:46 pm PST
Subject: Re: Stability of remote effects

[Martin Taylor 920213 19:00]
(Bill Powers 9202920213.1300)
>
>RE: getting thirsty
>
>This concept doesn't fit my stereotype of behaviorism. I had thought that
>behaviorists eschewed references to unobservable inner states.

Well, Jim taylor didn't, for sure, and he considered himself a classical behaviorist. But since Wayne informed me that dinosaurs are not extinct, I can't say.

Anyway, I wasn't talking about behaviourist, but about the behaviouristic method and its possibilities once one accepts the PCT model. The only relevant item in your posting about that is the question about How would one work back from stable environmental construct (the car staying in the middle of the road) to the ECS (or group of ECSs) responsible for that stability. I asked the same question about how does one work forward to the same construct. Bruce Nevin answered (correctly, I think) that it is a matter of empathetic intuition. What would I be controlling for if I were in this situation? It takes a more or less inspired guess to find a candidate for a controlled perception, but less work to check that it is controlled.

>
>Martin, you probably feel that you turned on a faucet and got hit with a
>firehose. I apologize for any overkill.
>
It was probably a good reference tutorial for new readers. I don't mind, but I think it was tangential to my point.
>
>
>Is all of this common knowledge among behaviorists?
>
I don't know, and I don't care. It could be, and they would be the better for it, but I doubt that it is.

Martin

Date: Thu Feb 13, 1992 6:27 pm PST
Subject: Diver

To: Bruce Nevin and other language lovers
From: David Goldstein
Subject: a book on Diver's approach
Date: 02/13/92

A while ago, I promised that as soon as I found a book reference on William Diver's approach to language, I would let you know.

Here it is:

Reid, W. (1991). Verb and Noun Number in English: A Functional Explanation. New York: Longman, Inc.

The book can be ordered by calling: 1-800-447-2226. The cost is $26.95 plus postage. The ISBN is: 0-582-29158-5 pbk.

I have not had a chance to read it. But I would like to share with you the following statements from page 40 which contrasts the Diver approach with other approaches(including Harris I would guess).

1. The view of language as a representational system is rejected in favour of a view of language as a communication instrument.

2. The sentence is replaced as the basic structual and analytical unit of language in favour of the linguistic sign.

3. The conception of language as an infinite set of sentences is replaced by the conception of a language as a finite set of signs.

4. The intersection of linguistic stucture and language use is shifted from a point between sentence and utterance, to one between linguistic sign and utterance.

5. The immediate object of explanation is the speaker's choice of linguistic signs in the communication of messages, not `sentence meaning' or `grammaticality'.

6. The staus of lexical and grammatical meanings as semantic components of the message is rejected in favour of the instrumental view of meanings as hints and clues.

7. Linguistic meanings are conceived as distinguishing tools rather than encoding tools.

8. Mapping relations between lingusistic meanings and messages are replaced by means-ends relations.

9. The relation of truth between linguistic meanings and messages is replaced by a relation of communicative efficacy.

10. The view of utterance formation as rule-governed is rejected in favour of a view of utterance formation as goal-directed.
11. The creative aspect of language is reassigned to its human users, rather than to formal characteristics of linguistic structure.

I believe that the Diver view of language makes a better marriage with the PCT view than that of Harris (or Chomsky). However, you can decide for yourself. It seems to come much closer to the sort of things which Bill Powers was saying about how an utterance is understood (through perceptions rather than through words).

Date: Fri Feb 14, 1992 4:04 am PST
Subject: bridges and bugs

From Greg Williams (9202014)

Rick Marken (910213)

I'm very happy to see that you seem to be making it through the storms OK!

>I don't recall seeing anything in the book or the article that could be called
>"matching his data to the neurophysiological data". What data does he want
>to match? Has this data been obtained from a real "working" bug.

Some of Beer's model is directly based on neurophysiological data. First, of a general sort (seen in many animals), are data on individual neuronal mechanisms (see discussions on pp. 25-34 and 49-53 of Beer's book), on synaptic mechanisms (including compound synapses, i.e., presynaptic inhibition; see pp. 35-37 and 56-58), and on simple neural circuits (see pp. 37-38). More importantly, data from cockroaches on gait circuitry (see pp. 71-74 and especially chapter 5, where Beer discusses his lesion experiments, which are analogous to certain types of neuroethological experiments) and edge following (see p. 120).

>But this level of match seems to be all that Beer is looking for -- for
>example, on p. 159 of the Discussion chapter he says "P. computatrix
>displays a number of interesting behavioral characteristics which are
>strickingly reminiscent of natural animals".

A behavioral match is not ALL Beer is looking for -- his additional constraint is neurophysiological reasonability at the neuronal level, as well as behavioral matching.

>I think the same could be said of the "Gatherings" program. And considerably
>more could be said for Bill's "Pigeon conditioning" program in which a
>simulated pigeon produces exactly the exact response rate data that is found
>in conditioning experiments (under many different schedules). Of course,
>Bill's model does not match the neurophysiology as exactly as Beer's
>presumably does.

Well, it might. But, with the current state of the art in neurophysiology, nobody can tell. A lot more is known about invertebrate neural circuitry than vertebrate neural circuitry, which makes modeling of invertebrates the better choice IF one wants those models to reflect the behavioral AND neurophysiological observations seen experimentally. There is certainly nothing wrong with using "global" (maybe the word is "lumped") models (PCT or otherwise) to attempt to predict and explain behavior of "higher" organisms. But if you go on to make detailed generative models even though there is insufficient data to make those models testable, then you will have a hard time with skeptics who are looking for correspondence with data, as well as internal coherence, consistency, and suggestiveness. They will claim (and have) that your models are really just metaphors. This is even a problem (though not in the end philosophically justifiable, I think) with "lumped" models -- some will say (and have) that they are just metaphorical. I don't think this is true in the case of "foundational" PCT models (the ECS as the overall model for behavior), but lacking generative details, some critics won't be swayed. My bigger point is that, generatively speaking, PCTers have tended to bite off more than they can support empirically (on the INSIDE) because of their organisms of choice. Ultimately, we all want to explain human psychology -- but right now, the details of generative models are testable to a significant extent only for "simple" organisms.

>But again, I have looked through Beer's book and I find absolutly no
>comparison of the physiological simulations in the model with physiological
>data obtained from actual behaving bugs.

Look again at the gait-circuitry material.

>Beer's ostensible goals are the same as the PCT goals (as you and Gary and
>others have pointed out) -- he really just wants to mimic BEHAVIOR.

Again, the additional constraint is "...using neuron-like components and experimentally supported circuitry." At least, that is the important goal, in my opinion, for PCT.

Regarding understanding how the gait circuitry works, Pat and I recommend carefully reading Beer's chapter on lesion studies and running NSCK with the L3R3 circuity -- modify stuff and see what happens! You'll learn far more, quicker, than from a verbal "explanation"!

Hope you continue to stay dry,

Greg
Date: Fri Feb 14, 1992 4:05 am PST
Subject: Avery Andrews remarks of 2/13/92

From: David Goldstein

Can you please explain: ..."it's in principle impossible and
pointless to try to predict anything about the properties
about a linguistic sign in advance--it's all post hoc
rationalization."f

I don't understand this.
Thanks.

****

Richard S. Marken   USMail: 10459 Holman Ave
The Aerospace Corporation   Los Angeles, CA 90024
Internet:marken@aerospace.aero.org
213 336-6214 (day)
213 474-0313 (evening)

Date: Fri Feb 14, 1992 10:38 am PST
Subject: Beahviorism; BEERBUG

[From Bill Powers (920214.1000)]

Martin Taylor (920213.1900) --

>I wasn't talking about behaviourist, but about the behaviouristic
>method and its possibilities once one accepts the PCT model.

What is the behavioristic method? Tutorial, please.

>The only relevant item in your posting about that is the question about
>How would one work back from stable environmental construct (the car
>staying in the middle of the road) to the ECS (or group of ECSs)
>responsible for that stability. I asked the same question about how does
>one work forward to the same construct.

Forward from the stable environmental construct? Forward from what? What do you mean by "forward to the same construct?" I'm getting confused.

Are you asking how one finds a model to explain the stability of the environmental construct? If that's what you're asking, I can lay out A method. It's possible to construct a CT model of any one control behavior based only on observable evidence, if that's what you're after. This model doesn't need a perceptual function, a comparator, or an output function -- it's all done in terms of visible relationships among visible variables. Before I launch myself into another tangential tutorial, I'd better know what point I'm supposed to be addressing.

Greg and Pat Williams (920214) --

Thanks for NSCK v. 3. I've printed out the listing for the first model, but as I haven't got Beer's book yet I can't visualize how it's put together. It's pretty hard to go from a listing of neural properties and the sublists of unexplained code names for the neurons that receive the output, to a structural diagram, and without a structural diagram I can't visualize how the thing is organized. I suppose I could try to draw the diagram from the lists of connections -- see you next month.

Thinking ahead to when I get those diagrams, I have some questions.

The basic one is, what are the constraints on building a model with these neural components? I assume there will be a list of sensors and the variables to which they respond (and their manner of response). There will be a basic neuron with inputs, outputs, and parameters that can be adjusted (or perhaps several neuron types). Also I suppose there will be a list of muscles and their attachments showing what they affect and what kind of effect they have. This much is fine -- it's basically a problem of understanding the properties of the components and their possible input-output relationships.

What I'm most concerned about are the connections. Do the neurophysiological data dictate which neuron receives inputs from which other neurons? Do the data specify what kinds of neurons must be connected to what sources and destinations? Are the signs and effects of the connections (excitatory, inhibitory, gating, etc) specified by neurological data? In other words, to what extent am I free to design a system, and towhat extent am I constrained to use connections and properties found in the real cockroach? I assume that since Beer's model is derived from neurophysiological data, there won't be a great deal of freedom here -- otherwise this would just be an exercise in circuit design using a somewhat unfamiliar parts kit.

If I had complete freedom I'd just design some control systems and an variable-frequency oscillator to drive the walking sequence. That would be (considering that I've spent 40 years designing electronic circuits with an ever-changing parts-bin) not too hard. But if there are rigid constraints set by data about the circuitry that actually exists, the design problem would be much more difficult (and significant).

Converting outputs to effects on the bug concerns me. When a leg is raised to move it back to an initial position, what determines its speed of movement? When the bug turns while some of its feet are in contact with the ground, how does the geometry work? Are the legs jointed in the middle or unjointed? When the body moves relative to the point of foot contact, a constant-length unjointed leg would go through some pretty odd swiveling, since it can't change length. Can this really be done with just two muscles per leg?

I can sort of visualize how the body might turn when supported by only two one-segment legs on each side -- turning and rocking vertically at the same time. But it seems to me that the computations would be rather horrendous, especially converting forces at the joint into effective forces on the body. How do you compute the body position during a turn? (Please don't refer me to the source code -- at this point I don't think I could find the relevant parts without some help). What are the degrees of freedom of the legs, and how are the legs affected in those degrees of freedom by the muscles? I guess that's the main question. I suspect that some cheating is going to be necessary here. But I don't want to re-invent the wheel.

Best to all,

Bill P.

Date: Fri Feb 14, 1992 1:51 pm PST
Subject: diverians

Chomskyans (& Harrissians, I believe) like to try to formulate rules that will predict what sentences are grammatical acceptable and what they mean, rules which in principle (and nowadays, sometimes in fact) can be implemented on computers. E.g. a sentence-structure rule might say something like:

```
   S   =  (XP:TOPIC)  NP:SUBJ  VP
```

That is, a sentence consists of an optional topic phrase, of any type, followed by a noun-phrase subject, followed by a verb-phrase. Diverians reject in principle any attempt to do this sort of thing.

They do, I might add, have some interesting ideas, which make sense in motivating language change. It's the whole package that I find unsatisfactory.

Avery.Andrews@anu.edu.au

Date: Fri Feb 14, 1992 3:54 pm PST
Subject: Re: Bill's ?? on bug details

From Pat & Greg Williams (920214)

>Bill Powers (920214.1000)

>The basic one is, what are the constraints on building a model with these
>neural components?
>What I'm most concerned about are the connections. Do the
>neurophysiological data dictate which neuron receives inputs from which
>other neurons? Do the data specify what kinds of neurons must be connected
>to what sources and destinations? Are the signs and effects of the
>connections (excitatory, inhibitory, gating, etc) specified by neurological
>data? In other words, to what extent am I free to design a system, and towhat
>extent am I constrained to use connections and properties found in the
>real cockroach?

PARTS (not all) of Beer's model are strongly constrained by neurophysiological data. You really need to refer to Beer's book (even if it costs $30), especially the parts Greg referred to in his most recent post to Rick Marken. The two unanswered questions we had for Dr. Beer after going through his book (and his answers) are these:

1. Q: When a foot is UP and the swing motor neuron for that foot's leg is producing a force, how does the motion of that leg relate to the force?
   A: (force/maximum force) * (maximum swing rate) = PI/15

2. Q: When a foot is down and a leg moves past its limits, what does "its angle will be restored to the proper range if it has bent too far forward, but not if it has bent too far back" [in the book] mean?
   A: It means that if bent too far forward, there is instantaneous springback to the forward limit.

All of your specific questions about geometric and physical constraints on leg positioning, etc., are answered in the book (mainly in Appendix A). Remember, this is a 2-dimensional model (Beer has since expanded into 3-d with his robot bug) and the physics is simplified to speed computations (F does not equal M * a, etc.; Beer has since expanded into real-world physics with his robot bug).

>I can sort of visualize how the body might turn when supported by only two
>one-segment legs on each side -- turning and rocking vertically at the same
>time. But it seems to me that the computations would be rather horrendous,
>especially converting forces at the joint into effective forces on the
>body. How do you compute the body position during a turn? (Please don't
>refer me to the source code -- at this point I don't think I could find the
>relevant parts without some help). What are the degrees of freedom of the
>legs, and how are the legs affected in those degrees of freedom by the
>muscles? I guess that's the main question. I suspect that some cheating is
>going to be necessary here. But I don't want to re-invent the wheel.

All of these questions are answered in the book, but it actually might be easier to look at the part of the source code that moves the bug. Expand the source code from SNS87.EXE (just run it). Then look at UPDATE.C. This is the heart of the source code where the neuron parameters and the bug state is updated. Somewhere in the middle is a comment that states: /* calculate bug position, etc. */. Look at the code after that, and you will see the very simple calculations that go into determining the forces and bug position. It's not very well commented, but I don't think it is all that unclear. Don't look (at least in the beginning) at the non-8087 versions of the code because they are uglified by all kinds of scaling and conversion that make it hard to understand.

Pat & Greg

Date: Fri Feb 14, 1992 4:22 pm PST
Subject: bridges and bugs

[From Rick Marken (920214)]

Greg Williams (9202014) says:

>Regarding understanding how the gait circuitry works, Pat and I recommend
>carefully reading Beer's chapter on lesion studies and running NSCK with the
>L3R3 circuity -- modify stuff and see what happens! You'll learn far more,
>quicker, than from a verbal "explanation"!

Can I ask some quick questions? The answers may be in the book but you can probably give them to me sooner.

1. There are are two sensors in each leg involved in "gait control"; the backward and the forward angle sensors. They seem to be described as binary devices. The forward angle sensor "encourages" the pacemakes (P) to terminate its current burst when the leg is "all the way forward". So it sounds less like an angle sensor than an event detector. Is this true? If the sensor responds only when the limb is at the limit of its throw, I don't see how the bug could control leg swing against disturbances that leave the leg with the range of its throw. Do real bugs have sensors like the one's in Beer's bug?

2. The foot output -- is it a continuous variable or binary (up/down)?

3. What is the function relating output to input variables? Beer doesn't say much about this. I presume it is linear-- eg -- the output of the forward angle sensor is proportional to the "swing" output (if the sensor is continuous -- if it is binary then the forward angle sensor output is 0 for every part of the swing except when the leg gets "all the way" forward). Is that right?

4. How does turning disturb the movements of the other legs? What happens to, say, L3 (one of the non-turning legs) when a turn is being executed. Maybe I could find this in the book but if you could tell me real quick it would help. The bug's compensation for the disturbances created by turning is, to me, the most interesting feature of the bug's behavior -- it's a sure sign that one of the perceptual variables (forward or backward angle or both) is being controlled. This may be a way to build a little bridge to PCT. If we can show what the legs are controlling then maybe we can predict the results of his ablation studies for him (so I won't look at those results until I figure out the equivalent control model for the bug).

5. Unfortunately, there seems to be no way to introduce external

disturbances to various aspects of the bugs gait. I would like to
try to "trip" it by preventing an elevated foot from being lowered for
stance phase to see how the bug compensates (if it does). I bet real
bugs do. Is it possible to introduce any kinds of external disturbance?
How about changes in the feedback function (relating outputs to inputs)?
These capabilities would make it possible to test for controlled variables.

Building, building, building.

Love (for Valentine's day)

Rick

Date: Fri Feb 14, 1992 4:23 pm PST
Subject: Re: Beahviorism; BEERBUG

[Martin Taylor 920214 17:40]
(Bill Powers 920214.1000)
>
>>I wasn't talking about behaviourist, but about the behaviouristic
>>method and its possibilities once one accepts the PCT model.
>
>What is the behavioristic method? Tutorial, please.

I am surely not the person to provide a tutorial on behavioristic methods, if
by that one means what the dinosaurs do. I'll try to say what I mean by it,
and if that disagrees with other people's descriptions, so be it.

Let's try this without reference to PCT, which is hard for me because the PCT
concepts have got so ingrained that it is difficult to see things in any
other way. But let's try.

There are three separable systems to worry about. E, the experimenter, who is
able to sense events in W, the world, one part of which is S, the subject.
E sees things happening in W, some of which are involve S. All of
these things that happen are initially incomprehensible. They just happen.

After some time, E determines that there are some apparent regularities in
the things that happen in W. E also can (in E's opinion) cause things to
happen in W, but for the time being let's ignore that. We will initially
develop an observational science. E sees that sometimes a pattern of events
in W will precede another pattern of events involving S (and at this stage,
"involving S" probably means only movements of S and of things moved by S
and perhaps things moved directly by things moved by S). Other patterns
of events in W will tend to occur after particular patterns of events
involving S (S moves a little lever on the wall and light appears).
All of this is still mysterious, but perhaps a little bit comprehensible.
Statistically, there is less information to be obtained by watching S
interact with W than the sum of the information obtained by watching S
and W seaprately. Aha! S is doing something with W! Can we see what?

It is raining. S appears with an umbrella. Sunny, no umbrella. And again
(usually, but not always). OK. Problem solved: Stimulus=rain, Response=
umbrella. Cause and effect. Aha! Got it! Publish! But why some days no
umbrella when it is raining, and some days sunny but still an umbrella? Bad
stimulus conditions? Bad response measure to correct stimulus conditions?
Or -- INSIGHT-- is something else happening?

Let's experiment. Hide the umbrella. Look some more. Now what do we see? Sunny days, S looks much as before, but is never seen with umbrella. Rainy days, still no umbrella. But what's this? Now on rainy days, S wears a hat and a long coat, never seen before except when S did not carry an umbrella. So what now is the response? R = (Carry umbrella or wear hat and coat)? Getting a bit complex, aren't we? But it must be true, because we see it.

Experiment more: hide hat. Now on rainy days S carries a newspaper over head. Oh--Response not properly defined before, perhaps. R = (umbrella or hat or newspaper)? Or -- INSIGHT -- response is "don't get wet". But is this a response? Have we found a drive, an intention, a gefarblworgle? S does what is necessary not to get wet when it rains. More experiments change how S looks, but does not change S staying dry.

At this point E has developed an internal abstraction for S, an intention to stay dry, or a drive to dryness, or some such. Does E notice that the behaviour of wearing or carrying protection depends on the perception of raininess by S. Probably so. We could still be in S-R mode. But E might notice that the choice of clothing depends not so much on the rain as on the difference between the probable wetness to be perceived if no protective clothing is chosen and the desired perception of dryness. If E notices that, which can be found more readily if E can do something that seems to affect Ss "dryness drive", then E is well on the way to discovering control, reference levels and the control hierarchy.

Of course, if E starts with the idea of a control hierarchy, then it is easier to invent experiments to determine what might be being controlled, which in turn means that it is easier to discover stable constructs that relate the W events to hypothesized abstractions in S.
>
>>The only relevant item in your posting about that is the question about
>>How would one work back from stable environmental construct (the car
>>staying in the middle of the road) to the ECS (or group of ECSs)
>>responsible for that stability. I asked the same question about how does
>>one work forward to the same construct.
>
>Forward from the stable environmental construct? Forward from what? What do
>you mean by "forward to the same construct?" I'm getting confused.

I'm not surprised you were getting confused. I meant forward "from" the same construct, by looking at what information might be available to S through S's sensors, as opposed to looking at what S does with S's muscles.
>
>Are you asking how one finds a model to explain the stability of the
>environmental construct? If that's what you're asking, I can lay out A
>method. It's possible to construct a CT model of any one control behavior
>based only on observable evidence, if that's what you're after. This model
>doesn't need a perceptual function, a comparator, or an output function --
>it's all done in terms of visible relationships among visible variables.

A tutorial on THAT would be worth having. If you really can do that, then I am astonished that you haven't become the GOM of experimental psychology. That's Newton stuff. Go to it.

>Before I launch myself into another tangential tutorial, I'd better know
>what point I'm supposed to be addressing.

None of the above, based on the previous interchanges. The point is whether the stability of the path from the controlled state in the world to the controller construct in the subject is greater if you try to determine what the subject can detect, discriminate, perceive, etc., or if you try to determine what he has moved, manipulated, constructed, etc. So far, I have no real opinion on the matter, other than that 100 or so degrees of freedom in join movements seems easier to handle than millions of degrees of freedom of sensor signals. But that could be a false notion; millions of degrees of freedom in a molecular gas are easier to handle than hundreds, or even tens.

Martin

Date: Fri Feb 14, 1992 6:34 pm PST
Subject: Beer bug details for Rick

From Pat & Greg Williams (920214)

>Rick Marken (920214)

>1. There are are two sensors in each leg involved in "gait control"; the
>backward and the forward angle sensors. They seem to be described as binary
>devices. The forward angle sensor "encourages" the pacemakes (P) to terminate
>its current burst when the leg is "all the way forward". So it sounds less
>like an angle sensor than an event detector. Is this true? If the sensor
>responds only when the limb is at the limit of its throw, I don't see how
>the bug could control leg swing against disturbances that leave the leg with
>the range of its throw. Do real bugs have sensors like the one's in Beer's bug?

The forward and backward angle sensors are either on or off. They can't control leg swing against disturbances that leave the leg not too far forward or backward. What effectively happens though usually is if the other legs are applying force, the leg will (if it is down) eventually hit one of the limits. If the foot is up the leg is swinging and will again hit a limit. On page 74 is a diagram showing such sensors in real bugs.

>2. The foot output -- is it a continuous variable or binary (up/down)?

It is binary. The details are in Appendix B.

>3. What is the function relating output to input variables? Beer doesn't
>say much about this. I presume it is linear-- eg -- the output of the
>forward angle sensor is proportional to the "swing" output (if the sensor
>is continuous -- if it is binary then the forward angle sensor output is 0
>for every part of the swing except when the leg gets "all the way" forward).
>Is that right?

As shown in Appendix B, the latter (binary) is correct.

>4. How does turning disturb the movements of the other legs? What happens
>to, say, L3 (one of the non-turning legs) when a turn is being executed. Maybe
>I could find this in the book but if you could tell me real quick it would
>help. The bug's compensation for the disturbances created by turning is,
>to me, the most interesting feature of the bug's behavior -- it's a sure
>sign that one of the perceptual variables (forward or backward angle or both)
>is being controlled. This may be a way to build a little bridge to PCT. If
>we can show what the legs are controlling then maybe we can predict the
>results of his ablation studies for him (so I won't look at those results
>until I figure out the equivalent control model for the bug).

If the foot of a leg is down, it stays where it is and the leg stretches to accomodate the bug movement; (not in book:) if it stretches past its forward limit then it instantaneously springs back to that limit, but it can stretch backwards indefinitely. If the foot is up, it isn't affected by the bug movement (it does what its neural inputs say). Note (not in book): When a leg is up, (force/maximum force) * (max swing rate) = PI/15. The two "not in book" facts were the only two details we had to call Dr. Beer about to complete the program. There are also a few typos in Appendix B.

>5. Unfortunately, there seems to be no way to introduce external
>disturbances to various aspects of the bugs gait. I would like to
>try to "trip" it by preventing an elevated foot from being lowered for
>stance phase to see how the bug compensates (if it does). I bet real
>bugs do. Is it possible to introduce any kinds of external disturbance?
>How about changes in the feedback function (relating outputs to inputs)?
>These capabilities would make it possible to test for controlled variables.

You can do anything if you reprogram it! But you can't do these sorts of things without changing the code. One thing you CAN do without reprogramming is experiment with injecting current into the LCS neuron (best done in the L3R3 setup) to vary walking speed and see different gait patterns. And you could do lesioning by setting some neural connections to zero ("dead" sensors or muscles).

Pat & Greg

9202C CSGnet

Date: Sat Feb 15, 1992 11:56 am PST
Subject: Language; BEERBUG; CT without model

[From Bill Powers (920215.0900)]

Avery Andrews (920214) --

You mention
>rules which in principle (and nowadays, sometimes in fact) can be
>implemented on computers. E.g. a sentence-structure rule might say
>something like:

>   S   = (XP:TOPIC) NP:SUBJ VP

How does such a computer program identify (perceive), in a source sentence, things such as XP, NP, and VP? This might tell us something about the human perceptions involved (unless these computers are programming themselves).

Pat & Greg Williams (920214) --

Thanks for clarifications. I know that some simplifications are always necessary in modeling, so don't think I'm faulting anyone for that. The simplifications do make it harder to construct a control-system model of the cockroach, because the feedback link determines how the inputs will behave under influence of the outputs. If that part of the loop isn't reasonably realistic, the behavior won't be, either. I may have to alter some assumptions, particularly the on-off nature of the sensors (unless you can tell me that they did experiments with the roaches' legs, measuring sensory signals, and discovered that the angle sensors really are on-off). I rather doubt that these signals are on-off, because when the roach pauses, some of its legs could be in an intermediate position, and the roach wouldn't know which way to move them when walking began again.

Just fiddling, I changed the SWL1 motor neuron capacitance to 10,000 units, so its voltage output changed VERY slowly compared to the other neurons. I couldn't see any obvious difference in the behavior of that leg, except perhaps it moved a little out of phase with its opposite leg.

It looks as if all the PLx neurons are connected in a ring, in a mutually-inhibitory mode that creates positive-feedback oscillations (or would in a continuous system). These seem to be comparators, but with lateral relationships between them at the same level. This is the arrangement I thought I remembered, which appears also in other neural circuit diagrams I've seen. The result would be like a wave propagating down each side of the bug at the PLx level, creating the walking sequence (except for lifting the legs). It will be interesting to convert all these binary neurons into continuous-variable neurons, as I presume can be done by adjusting gains and integration constants. If the PLx neurons are made continuous, the wave-like propagation of effects will provide effective reference signals for the leg-position systems that are sine-waves in an out-of-phase relationship. The legs won't necessarily move in sine-waves, because during part of each wave the feet are stuck to the floor and the leg forces will interact through the body.

Martin Taylor (920214.1740) --

Nice exposition on the behaviorist method. It sounds a lot like science.
>At this point E has developed an internal abstraction for S, an >intention to stay dry, or a drive to dryness, or some such.

This is the first step in the test for the controlled variable, although a better guess at the reference level is possible: an intention to keep wetness at zero. It's hard to characterize the effect of rain on dryness, so the variable involved really should be wetness. But then the desired condition is in different units from the condition as perceived. So convert the desired condition to the same units: zero wetness. We want to express the reference level as one possible state of the variable. The variable is wetness; its reference level is apparently zero, unless you like getting wet.

>E is well on the way to discovering control, reference levels and the control hierarchy.

Correct. Here's how to go from there to identifying a control system without a model.

First, characterize observables in the environment. Note the degree of wetness W. Note effect A of action on same variable. Note effect R of rain-rate on same variable (measure both rain and action in terms of effect on wetness). Note that $W = R - A$. Equation says that if $R = 0$ and $A = 0$, $W = 0$. Also predict that if $R > 0$ and $A = 0$, $W > 0$. Prediction verified by observation: so far so good. If $R = 0$ and $A > 0$, $W < 0$. Prediction no good -- can't have negative wetness. So equation applies only when $W >= 0$.

Now postulate that $A = g(W - W^*)$. Amount of action is based on degree of wetness in relationship to some base wetness, W^*, to be determined from data. This leads from $W = R - A$ to

$W = R - g(W - W^*)$. g is the form of the unobserved organism function.

Solve W = R - g(W - W*) for the form of function g that will satisfy the observed relationship between data points for R and W. W* is the value of W at which action has zero effect on wetness. The constant W* is the reference level (not the reference signal, which would be a variable inside the actor). We find a function g that will fit W = R - g(W - W*) to the recorded values of W and R, obtaining W* from the data.

The function g describes the effect of stimulus W, via the actor, on action A. Calculate the partial derivative g(W-W*), the slope of the stimulus effect on behavior, and the partial of W with respect to A, the slope of the effect of behavior on the stimulus (which is observable). If the product of these two partials is a large number, and if that product is negative, we have a control system. The ideal control system of that form, in which this loop gain is negative infinity, predicts that

W = W* (wetness equals reference level of wetness)   A = -R (effect of action on wetness equal and opposite to effect of rain on wetness)

Notice: no input function, comparator, or output function.
All variables observable.
The only postulate is that the organism is reponding to W by producing A, so that A is some organism-function of W. If this is false, the best form of g found from the data will be

A = 0 * (W - W*)

-- in other words, wetness depends on the organism's action and the rain-rate, but the action does not depend on wetness. The appearance of stimulus and response is due to statistical happenstance.

If a systematic form of g is found, then we may or may not have a control system. Only if the loop gain (product of the partials) is negative and large is control of W relative to W* verified.

The above just shows the logical form of the procedure. The actual measurements and mathematical relationships you use would be selected as appropriate. For instanced, wetness might actually depend on rain rate as a time integral minus an exponential decay (evaporation), and so on.

This analysis is important for any empirical approach to behavior. When you see one variable depending on another, S-R fashion, the form of the observed relationship is NOT in general a description of the organism's response to the stimulus. To find the true form of the organism's response function, you have to examine all links connecting the observed variables, input, output, and independent, in both directions. Only by solving the resulting simultaneous equations can you deduce the true organism function.

I think behavioral experimenters have sensed this problem. This is why they have defined stimuli and responses mainly as events taking place at separate moments. That allows considering external events separately and independently of the response events. In this way behavior can be represented as an alternating sequence: S->R, then R-> next S. Of course the answers you get from such an analysis are wrong.

Date: Sat Feb 15, 1992 12:37 pm PST
Subject: Real bug leg position sensors

From Pat & Greg Williams (920215)

Bill Powers asked about cockroaches' leg position sensors. In Dr. Beer's book, he references work by K.G. Pearson and colleagues (see p. 73 and diagram on p. 74 of INTELLIGENCE AS ADAPTIVE BEHAVIOR) and glosses that work as follows (p. 73):

"The central pattern generator [found by Pearson, et al.] underlying the rhythmic movements of each leg is shaped by feedback from two sensory structures. Tiny hairs near the leg joints which are stimulated when the leg has reached its extreme forward position have been shown to inhibit the CPG and the swing motor neurons and excite the motor neurons controlling stance. These hairs appear to play a role in controlling the switch from swing to stance. [sounds basically on-off]
"Sensory structures which measure stress in the legs have also been shown to influence locomotion in two ways. First, their stimulation can prevent activity in the pattern generator, and thus may play a role in controlling the switch from stance to swing by preventing a heavily loaded leg from swinging. Second, these sensory structures excite stance motor neurons, providing a possible mechanism for load compensation." [continuous?]

The gait circuit diagram on p. 74 (from Pearson) shows the cuticle-stress receptors excited by stance, and the hair receptors excited by swing, but no info on either path's on-off/continuous nature.

On p. 77, Beer says: "When a leg is all the way back, the BACKWARD ANGLE SENSOR encourages P to initiate a swing by exciting it. This sensor is related to the stress receptors in Pearson's model. However, the correspondence is not nearly as direct as for the forward angle sensor. The stress receptors are essentially load sensors, while the backward angle sensor is a position sensor. A variety of load sensors were experimented with, but with largely unsatisfactory results...."

Why not contact Dr. Beer on some of the technical details. We think he would be happy to know you are down in the thick of the circuitry!

Hope this helps,

Pat & Greg

Date: Sat Feb 15, 1992 6:23 pm PST
Subject: language (grammar rules)

There are lots of ways in which the subordinate phrase types in a rule such as:

```
  S  ->  (XP:TOPIC) NP:SUBJ  VP
```

can be recognized. E.g., in `pandemonium'-type recognizers there would be a circuit attempting to recognize each phrase-type, these circuits calling each other, & the whole thing bottoming out in a lexicon specifying the phrase-type of individual words.

Pandemonium seems to me to have problems with recursion, but maybe there's a way around this. At any rate, I'm about to put my toy LFG system on an ftp server, & will make an announcement when I've done this, so people can play with it.

Avery.Andrews@anu.edu.au

Date: Sun Feb 16, 1992 12:32 pm PST
Subject: Behavioral test

[From Rick Marken (920216)]

Bill Powers (920215.0900) in reply to Martin Taylor gives the details of a "behavioral" approach to testing for controlled variables.

Sometimes it's almost TOO good.

What a lovely post, Bill. Suitable for framing.

Thanks

Rick

Date: Sun Feb 16, 1992 3:53 pm PST
Subject: BOURBON ON BEER

[From Tom Bourbon -- 920216 17:40]

During the past couple of months, while most of the *sturm und drang* of the debate on BEERBUGS unfolded, I could only watch, due to hardware, software and administrative problems with access to the nets from our campus. From time to time, I could read the mail that accumulated during a few days, or a few weeks, when there was no access, or I could discover that all mail during that time bounced off of our system. Silently watching fragments of the debate and re-reading most of Beer's book, left me with the following thoughts.

I. BEER'S WORK. His modeling and simulations are a major feat of software engineering. The model is complex and it simulates in real time. I am impressed. But by "software engineering" I also mean that he designs software "components" for *Periplaneta computatrix* -- components that reflect a design strategy in which a desired end is identified (a "behavior" that is interpreted as a completed action, an end result that requires direct explanation). The result is a "neural" function that produces the specific end identified at the start of the design process. The finished system behaves in strict accordance with the principles designed into it -- it creates behavior as an end result, out of "commands" and "sensory signals."
II. PHYSIOLOGICAL REALISM? Pat and Greg Williams have argued that Beer's modeling is more significant than PCT modeling due to Beer's incorporation of *principles* and *data* from the neurosciences (or whatever name they used for the biology-behavior combine). However, the more I read Beer's book, the more I realize he uses CONJECTURES from the neurosciences, not facts that make his model inherently superior. His finished product is an hypothetical system (and an IMPRESSIVE one), with a PARTIAL grounding in hard data about nervous systems, whether in roaches, or other beasts. Many of the data and conjectures he uses come from work driven by the notion that behavior results from a C-->E process. (Beer is not "at fault"

on this; such is the state of knowledge in the field.)
A. Beer picks target "behaviors" and designs plausible component circuits to produce exactly those behaviors -- one at a time -- then he assembles the components from his software parts kit into a complex system. Most neuroscientists follow that process in reverse. They identify an isolated "behavior," as an accomplished event or process, then they look for a component or subsystem of a living system that seems to produce approximately the same behavior. In living systems, stimulation studies and lesion studies are most often aimed at identifying antecedent causes of behavior. A frequent driving assumption is that "commands" act WITHIN a nervous system, and "inputs" act ON it, with the net effect of the two determining behavior. Beer's descriptions of the work he cites throughout the book seem to me to embody that concept.

B. Independently of the issues in "A," Beer frequently reports (correctly) that there are no data on the neurological substrates of some of the behaviors he wants to model, and he often says that many of his components are derived from conjectures in the neurosciences, not from data. He incorporates features and interactions that he and others BELIEVE would explain something, or that APPEAR TO, MAY, of MIGHT explain something, or that are PROPOSED explanations of something. I do not say this to be critical of Beer, but anyone who feels threatened due to a belief that Beer argues from hard data in physiology, exclusively, need have no such fears. Data of that kind are not available.
END OF PART 1. I am typing this while online on a machine that is "off limits" to faculty. I could be bumped at any time , so I will post this, then start on PART 2, a selection of several quotes from Beer that illustrate my points.

Tom Bourbon   <TBourbon@SFAustin.BitNet>

Date: Sun Feb 16, 1992 5:21 pm PST
Subject: BOURBON ON BEER: PT 2

[From Tom Bourbon -- 920216 19:15]
This is a continuation of a post at 17:40. Sure enough, I was bumped 5/6th of the way through my first try at this continuation. This will become the second of three parts.

A FEW EXAMPLES OF BEER'S PHYSIOLOGICAL REALISM

1. The pacemakers that drive the legs of Beer's model are an instantiation of a previously unsimulated model described by Pearson in the 1970s. Apparently Pearson studied bejavior and nervous systems in real roaches and summarized his obserevations in a descriptive model that Beer reproduces in his book. Beer frequently refers to Pearson's model, then talks about neurons IN THE MODEL. This is fine, but readers must be alert to the fact that Beer is tallking about MODELED NEURONS, not real ones. The activity of those pacemakers, alone and in coordination, is the heart of Beer's model. They comprise the "central pattern generator" that creates many of the outwardly correct features of movement by the model. Om p. 73, Beer reports that, "Whie the neural circuitry comprising this pattern generator REMAINS UNKNOWN, nonspiking interneurons have been identi-

fied which APPEAR TO BE members of this pattern generator." (My emphases.) Also on p. 73, he reports on some physiological data. "This pattern generator operates against a background of steady excitation descending from higher brain centers, which also excites the pattern generator itself." What is a "steady background of excitation" and how does it differ from the steady excitation that acts on the generator itself? This report of descending excitation and of its role is not a pure DESCRIPTION, if any such thing is possible; it is an interpretation of neural functioning. Could that background ne the collective set of reference signals from higher brain centers? There is no way to know, since most research on physiological substrates of behavior occurs in settings designed to identify linear causal mechanisms: If control is present, it is often masked, or eliminated by "experimental controls." Greg, are you familiar with the data to which Beer refers, apparently after Pearson?

2. In his discussions of how the activity of the CPG is shaped by feedback from sensory structures, Beer reports that the hair cells located at the maximum forward extent of the legs "APPEAR TO play a role" in, and that stress sensors "provide a POSSIBLE mechanism for," feedback. (My emphases.) He models the structures as though they play those roles, but apparently the original data are not clear, or convincing.

3. Each leg has a pacemaker. The mechanism by which the six pacemakers become synchronized or coordinated is obviously a significant feature of the behavior of real bugs and of the model. The explanation used by Beer, apparently after Pearson, begins with coordination as an accomplished fact, then seeks a mechanism that will produce exactly that coordination. Pat and Greg, this is one reason I lean toward questioning the claim that Beer's is NECESSARILY a generative model -- but I know some additional coordinative phenomena do emerge out of his mechanisms.

In his description of the mechanism for coordination, Beer writes (p. 78), "In order to ENFORCE the generation of metachronal waves (TB: one form of coordination), the following IDEA from Graham was used ...." (My emphases.) And on p. 94 he writes that IN THE MODEL, "The connections between pacemakers ARE DESIGNED TO provide decentralized coordination of their relative phases." (My emphases.) This is a description of a useful tactic on modeling, not of a fact of biology. Small wonder that a system designed to produce a specific type of coordination produces exactly that type of coordination. (In no way do I imply that he task was easy or that the result is trivial. All I mean is that this is a nice feat of engineering, not a description of established neuroscience.) END OF PART 2.

THE COMPUTER COPS STILL HAVEN'T CAUGHT ME. PART 3, THE LAST, FOLLOWS

Tom Bourbon <TBourbon@SFAustin.BitNet>

Date: Sun Feb 16, 1992 5:54 pm PST
Subject: BOURBON ON BEER: PT 3

[From Tom Bourbon -- 920216 19:45]

Part three of a post at 17:40 and 19:15.

FINAL EXAMPLES OF PHYSIOLOGICAL REALISM IN BEER'S BUG

4. p.86."The relationship between central and peripheral components in insect locomotion in particular and pattern generation in general is currently the topic of some controversy in neurobiology." (There is no factual solution, it is a matter of theory. PCT, anyone?)

5. p. 87. "... A considerably deeper understanding of the pattern generation circuitry responsible for locomotion in the cockroach must be atreined before a significantly more realistic model could be constructed." (Might it matter were the people who studied real roaches in search of that deeper understanding to test the hypothesis that the nervous system of the roach serves the behavioral process of control?)

6. What most endears the BEER BUG to many people is the way it turns, wanders, recoils from barriers, and follows edges. Beer categorizes all of those "behaviors" as "exploration." Consider this, from p. 107: "Because NO neural circuitry for ANY exploratory behavior is currently available, the controllers in this chapter were DESIGNED FROM SCRATCH. However, they make use of intrinsic properties and architectures which are SIMILAR IN SPIRIT to those found in natural animals." The natural animals to which Beer refers are any kinds from which data seem useful, not just roaches. We all probably know what Beer means by "similar in spirit," but what he describes is a way of building a model, not a firm appeal to hard biology. This is not a criticism of Beer, just a reminder that no one need feel intimidated by frequent allusions to biology. The PCT architecture is also "similar in spirit" to known principles and architectures: positive (excitatory) processes meet negative (inhibitory) ones at a comparator (cell body) where their interaction determines the influences (error signals) that act on the next components. Beer is no more in touch with biology than is PCT, at that level of detail.

7. When Beer builds his circuits for feeding behavior, he argues that positive and negative feedback are involved, but the evidence for positive feedback is weak -- of the sort that PCT people have often had to counter in the past. On that point, Beer is presenting an interpretation of the nervous system, not a factual account.
BRIEF CONCLUSIONS. Close reading of Beer's book convinces me that this is not the sort of thing you will find inkindergarten! It is an impressive piece of engineering. But it is not an exposition of well-established neurobiological facts. And I disagree with Pat and Greg when they imply (or declare?) that any measure of superiority accrues to a theory or model when its author alludes to hypothetical anatomy or to conjectural physiology. Those allusions are not of necessity wrong, but neither are they of necessity correct. And lest someone take offense, I do not hold Beer guilty of the common practices of neuroscientists -- unless he engages in those practices along with the others.

For my part, I hope to see BEERBUGGERS and STICKWIGGLERS thrash it out (mano-a-mano, or network-a-network) at the level of detail recently reached in the posts on CSG-L. Both of our houses have an adequate supply of issues left unadressed, or only partially addressed. Both groups work with the CNS, in the sense often described by the late Donald Hebb -- the Conceptual Nervous System. There is more than one CNS. Is one "superior?" Only time and successful modeling will tell.

Tom Bourbon <TBourbon@SFAustin.BitNet>
Dept. of Psychology
Stephen F. Austin State Univ.
Nacogdoches, TX 75962 Ph. (409)568-4402

Date: Sun Feb 16, 1992 8:21 pm PST
Subject: Re: Bourbon on Beer

From Greg Williams (920217)

>However, the more I read Beer's book, the more I realize he uses CONJECTURES
>from the neurosciences, not facts that make his model inherently superior.
>His finished product is an hypothetical system (and an IMPRESSIVE one), with a
>PARTIAL grounding in hard data about nervous systems, whether in roaches, or
>other beasts.

I think that both the facts (there ARE some) and the conjectures (yes, several of those, also) upon which Beer bases his model do make his bug model not necessarily "inherently superior" to Bill's detailed hierarchical PCT model, but more plausible and better grounded. To an extent, one's opinion on this probably reflects one's vested interests, but I bet that a majority of neurophysiologists would agree with me. As I said before, Bill's model MIGHT BE RIGHT -- but there is now, and will be for some time, much less data to support or contradict (any) models of the human control structure than to support or contradict models of "simple" organisms' control structures. In particular, there is evidence from neuroethological studies of, for example, nematodes, in support of the "inhibition-control" technique employed by Beer's bug to avoid conflicts among behavioral "modes," and in contradiction to Bill's reference-level-resetting technique. Yes, the facts and conjectures upon which the bug model is based might be skewed by S-R thinking. Still, Bill's "functional"-level modeling of humans is a good deal more distant from the data than is Beer's "neural"-level modeling of bugs. ALL models are hypothetical, but some are more so than others. With regard to the PCT "agenda," as Mary calls it, I think the fact that you and I disagree about the relative physiological grounding of Beer's and Bill's models isn't as important as the apparent fact that modeling ala Beer is generating a lot more interest than modeling ala Bill. I suggest that the implicit difference in both actual and potential empirical support for the two endeavors is part (but certainly not all) of the reason.

Many of the data and conjectures he uses come from work driven by the notion that behavior results from a C-->E process. (Beer is not "at fault" on this; such is the state of knowledge in the field.) A.

>Beer picks target "behaviors" and designs plausible component circuits to
>produce exactly those behaviors -- one at a time --then he assembles the
>components from his software parts kit into a complex system.

How does this differ from (say, Bill's arm-) modeling in PCT?

>He incorporates features and interactions that he and others BELIEVE would
>explain something, or that APPEAR TO, MAY, of MIGHT explain something, or that
>are PROPOSED explanations of something.

How does this differ from modeling done by control theorists? The only way I think it differs is that the "beliefs" used by Beer, in some cases, have a quite direct relationship to physiological data. Such a direct relationship never seems to exist for detailed PCT models.

>I do not say this to be critical of Beer, but anyone who feels threatened due
>to a belief that Beer argues from hard data in physiology, exclusively, need
>have no such fears. Data of that kind are not available.

I didn't mean to imply that Beer's model was constructed EXCLUSIVELY from "hard data in physiology." But it IS constructed with (in my opinion) reasonably realistic INDIVIDUAL-neuron models, regardless of your qualms about the worth of the physiologists' data and models for neuronal CIRCUITS of real cockroaches. Working at the "neuronal" level in real time is a significant accomplishment, as you have noted. Here again, I suggest that this accomplishment contrasts (as valued by behavioral scientists and roboticists, not just by me) with "functional"-level PCT models. Of course, control theorists can remedy this by making models at the "neuronal" level -- I think Bill and Rick are doing this.

>Greg, are you familiar with the data to which Beer refers, apparently after
>Pearson?

I have a lot of information on neuroethological studies on cockroach walking, mostly post-Pearson. If you are interested in the latest data, start with a search on Fred Delcomyn's papers.

>The explanation used by Beer, apparently after Pearson, begins with
>coordination as an accomplished fact, then seeks a mechanism that will produce
>exactly that coordination. Pat and Greg, this is one reason I lean toward
>questioning the claim that Beer's is NECESSARILY a generative model -- but I
>know some additional coordinative phenomena do emerge out of his mechanisms.

Beer DID seem genuinely surprised that the model mimicked gait transitions in real cockroaches so well!

>4. p.86."The relationship between central and peripheral components in
>insect locomotion in particular and pattern generation in general is
>currently the topic of some controversy in neurobiology." (There is
>no factual solution, it is a matter of theory. PCT, anyone?)

This and your other examples illustrate Beer's care in not overclaiming FULL empirical support for parts of his model. Beer is saying where he is conjecturing and to what an extent. I think this is admirable. Another point which I guess is obvious but I'll make anyway is that Beer went looking for the data -- but of course didn't find everything he would have liked. That too is admirable: to begin the modeling process by seeing what is known (and whether ANYTHING is known!) in detail about what you want to model.

>5. p. 87. "... A considerably deeper understanding of the pattern
>generation circuitry responsible for locomotion in the cockroach must
>be atreined before a significantly more realistic model could be
>constructed."

Apparently, the basic P.c. model -- maybe somewhat enhanced, I'm not sure -- works pretty well in Beer's 3-D robot.

>The PCT architecture is also "similar in spirit" to known principles and
>architectures: positive (excitatory) processes meet negative (inhibitory) ones

>at a comparator (cell body) where their interaction determines the influences
>(error signals) that act on the next components. Beer is no more in touch with
>biology than is PCT, at that level of detail.

I think you are minimizing the extent to which Beer relies on empirical evidence relative to the extent which Bill (in his detailed hierarchical model) relies on empirical evidence. Beer's cell bodies are pretty realistic, with leaky-integration dynamics, firing thresholds, saturations, etc. The temptation is great in "functional"-level models to ignore all that sort of stuff and put in a simple linear comparator.

>And I disagree with Pat and Greg when they imply (or declare?) that any
>measure of superiority accrues to a theory or model when its author alludes to
>hypothetical anatomy or to conjectural physiology.

The measure of superiority accruing to ANY model rests, in part, on its basis in TO-SOME-EXTENT-hypothetical physiology. Models are models, not facts. The question is how far from the available physiological data one's model is. I once met the (I think) Nobel Prize winner Alex Rich (I think he elucidated the structure of TRNA) and we discussed a book in which a fellow set forth in exhaustive detail how a series of "clocks" in DNA could regulate organismic development to the nth degree. Molecular biologist Rich deferred commenting on the worth of the book: "I'm not qualified. No data are available." Some (not ALL!) data are available regarding modeling the details of "simple" organisms' control structures -- vastly more than are available regarding modeling the details of human control structures. To leap to vastly unsupportable models of human control structures is to ask for peer review comments along the lines of Rich's.

>Both groups work with the CNS, in the sense often described by the late Donald
>Hebb -- the Conceptual Nervous System. There is more than one CNS. Is one
>"superior?" Only time and successful modeling will tell.

My overall point is that there are several organisms with CNSs to be modeled. At this time, I think it makes good tactical sense to model the simpler ones. Then you have more (but still not PERFECT) empirical support to fall back on when the reviewers ask. Regarding "superiority," I suspect -- with very little data to fall back on! -- that it isn't a question of either-or. Why couldn't there be various detailed control mechanisms in various animals, or even the same animal? Again I say why must PCTers cling to ONLY Bill's (detailed) model, given the paucity of data (but some negative data)?

A final question. How come this Beerbug debate gets framed as "threatening" to PCTers? I'M a PCTer, and I don't feel threatened. I think PCTers should see Beer's work as an OPPORTUNITY, not as a THREAT. Modeling at the bug (or worm, or other "simple" organism) level is an opportunity to (1) show that the overall ("lumped?") PCT model WORKS in an attention-getting way, (2) develop and experiment with empirically based detailed control mechanisms, and (3) basically get folks interested in PCT ideas. Who feels threatened, and by what?

Best,

Greg

Date: Sun Feb 16, 1992 9:51 pm PST
Subject: LFG system (grammar)

I've put my `LFG' system (for IBM PC compatibles) onto an ftp server, which may be accessed as follows:

anonymous ftp to durras@anu.edu.au

cd to pub/lfg

get or mget the files.

There are two compressed (.tar.Z) tar archives, plus their tables of contents. One, lfg-pc16.tar.Z is the `current' version of the system, while the other, lfg-pc_extras.tar.Z, contains an unfinished version with a nicer user interface, plus some selections from the source code (PDC Prolog; once Turbo Prolog).

So, what is LFG, anyway? It's one of the current `Chomskyan' grammatical theories (Chomskyan in a general sense - he is actually currently promoting some rather different), designed with an eye to making it possible to write linguistically decent grammars that are also computationally plausible.

It posits two levels of grammatical representation, `c-structure', which represents the overt structure of a sentence in terms of a hierarchy of phrase-types, and `f-structure', which represents the semantically significant grammatical relations, and is rather like a record in Pascal, or a struct in C, but much more loosely typed.

f-structures include something very much like the operator-argument structure that Bruce Nevin has occasionally talked about: most f-structures contains a `PRED' feature, that specifies a name of a concept, and the grammatical relations of its arguments, e.g., the f-structure of `the dog bit the cat' will contain the feature:

```
  PRED  `Bite(SUBJ,OBJ)'
```

The f-structures of `dog' and `cat' will appear in the whole f-structure as values of SUBJ and OBJ, and the tense value of the verb as value of TENSE, so the whole f-structure will be:

```
  PRED  `Bite(SUBJ,OBJ)'
  SUBJ   [PRED  `Dog']
  OBJ    [PRED  `Cat']
  TENSE  PAST
```

(linear order of the vertical lists doesn't matter).

The things in quotes should be thought of as abbreviations for, or pointers to, the definitions of concepts (which are presumably in part nonverbal for at least some of the vocabulary).

These f-structures are constructed by the joint operation of

a) annotated c-structure rules
b) lexical entries

The rules say things like:

```
   S  -> NP:SUBJ   VP.
   VP    V  (NP:OBJ)  (NP:OBJ).
```

i.e.: a sentence consists of an NP (with the grammatical relation subject) followed by a VP.
a VP consists of a verb, followed by an optional NP (with GR object) followed by another optional NP (with GR 2nd object (as in `John gave Sally money')

One can think of this as a scheme for supplying operators with their arguments, which is supposed to function under various perturbations of overt linguistic form that linguists think they know a fair amount about.

The package has been used in teaching, & there's a developing textbook to go with it. I don't know how much sense non-linguists will be able to make out of it on their own, but there it is anyway.

Avery.Andrews@anu.edu.au

Date: Mon Feb 17, 1992 10:44 am PST
Subject: BEERBUG

[From Bill Powers (920217.1000)]

Rick Marken (920216) --

Thanks.

Tom Bourbon (920216a,b,c); Greg & Pat Williams (920216, previous)--

Crime pays, Tom. Thanks for a very informative summary of the strengths and weaknesses of Beer's model. I agree that it is a considerable engineering achievement.

Greg & Pat: I would use neural data in my models if I could. I'm prevented in part by my lack of access to literature (in any easy way) and my priorities. I'm also prevented by many of the factors Tom mentions: lack of data that's better than an educated guess. The data that are available show many of the biases of neurological researchers -- the emphasis on thresholds, for example, and the use of on-off sensors and "state" output variables. Those are purely theoretical ideas, and they carry the flavor of belief that single impulses are significant and represent digital variables -- even though "neural currents" are used. The concept of continuous neural functions has been adapted to allow the use of digital variables after all!

There's some confusion in Beer's model, I think, between thresholds of firing and thresholds of response in terms of frequency inputs and outputs. Even if there are thresholds for generation of a single spike, there need be no thresholds when you measure output frequency as a function of input frequency (although there will be nonlinearity at the lowest input rates, and perhaps some threshold effect). Looking at the plots in your program, I see that the current outputs tend to jump from one value to another, although the voltage measures are smooth. I presume this is the threshold effect showing up, converting an analog relationship to a digital one by using high gain and a threshold -- much as TTL logic works. But this is a matter of how you adjust the parameters of the neuron, which I presume is not an empirically-based choice. Different choices would have led to an analog-computer model.

The ring-connected pattern generator is a good design idea even if it hasn't been traced yet in the cockroach. Mutual inhibition of neurons at the same organizational level will, as I've said, create positive feedback loops that will oscillate if the gain is high enough, and a ring or string of such units will create wave-like effects which will reflect as the gait of a multilegged creature. Something like this would be needed even in the human being at the fourth level, events, to create repetitive patterns of reference signals for lower-order systems (for walking and other kinds of temporally patterned packages of behavior). I'm not totally convinced that the "oscillator" concept is the only good one, however. Applied to human walking, it doesn't explain how walking can slow down, stop, reverse, go forward again, etc. It wouldn't do for a model of the way Michael Jordan moon-walks!

My model is empirically based, perhaps considerably more so than Beer's. "Empirical" doesn't necessarily imply dissection; it means only building explanations on observations. I chose to start with observed behavior, and to try to characterize organizations that will produce what we observe. Themain phenomenon we observe is control. As I showed in my last post, it's possible to deduce the presence of a control system strictly on the basis of externally-measurable variables, without a block diagram or even a nervous system. In a similar way it would be possible to show that a hierarchy of control exists, by observing cases in which a variable already characterized as a control phenomenon becomes the means of controlling another variable that is in part a function of the first one.

I have chosen to propose a simple functional model of the insides of the behaving system, based partly on my experience with electronic control systems, that would produce the kinds of control behavior we can see. I've said many times that this is a "canonical" model, meaning that it is a form that's equivalent in function to many other possible forms. In the brain stem, for example, the motor nuclei receive efferent signals from higher systems, and also "collateral" signals branching off from afferent signals at the same level. At this level, in every place in the brainstem I've seen pictured, the collateral signals have an excitatory effect in the motor nuclei, while the signals from higher systems have an inhibitory effect (an upside-down comparator). So there is no distinct comparison function at this level: the comparator is part of the output function. But the result is functionally equivalent to having a separate comparator with an error signal that enters an output function -- the mathematical description of the result would be the same. I have, as you can see, used such neurological data as I have found, although it would be futile to try to make the details any finer, given the state of neurological knowledge.

In the cockroach, it appears that the PLx neurons combine a perceptual and comparison function. The perceptual signals are not created by a distinct perceptual function where inputs are added and subtracted together. Instead, the inputs have additive and subtractive effects directly on the PLx neuron, which also receives a reference signal from an "LC" neuron. The effect, mathematically, is just as if the various perceptual signals reaching the PLx neuron were first received by an input function and combined mathematically to produce a distinct perceptual signal, with the result sent to a comparator also receiving the "LC" signal. The only difference from the perceptual part of my model is that this effective perceptual signal never appears as a single signal, and so no copy of it can be sent to higher systems. This is a very simple organism; on the face of it, higher systems make no use of this perceptual information. On the other hand, there is no evidence that such a perceptual signal does not exist, perhaps in the dendrites of the PLx neurons, being transmitted to higher systems through one of the kinds of synaptic connections not considered in Beer's model -- dendro-dendritic synapses, for example. And there is no assurance that the connections shown in the model are those that actually exist, according to Tom's report.

Another difference from my model -- an important structural difference from the canonical model -- is the cross-connection between systems at the same level. I've mentioned that possibility many times, but have never seen a good way to incorporate it into the model. Maybe Beer (or those whose work he used) has shown a good way. I will certainly try it out.

I agree with you that models which demonstrate interesting things, like walking bugs, capture attention, particularly if they contain at least a few control processes and show interesting behavior in relation to anenvironment. The implementation of such models using pseudo-neurons, however, is basically a trick, because only the mathematical relationships represented by those neurons actually matter -- until we can say that THIS is the true circuit diagram, THESE are the true parameters of the individual neurons, and THOSE are the correct models of the sensors, muscles, and physical consequences of action. Without having pinned these fundamental aspects of the circuit down, any model is simply a way of visualising relationships that are theoretically proposed, and the product is not an analysis of the real system's details, but a design that may or may not be the best one. Nature may not have settled on the best design, but we can't know which of the infinity of possible designs is actually used until we no longer have a choice about how to interconnect the neurons and how to adjust their parameters.

While we do still have a choice, it's much more realistic to just draw boxes and assign functions to them, looking for a design that will economically accomplish the behavior we observe from outside. How the function in each box is accomplished in neurons is a matter for investigation at another level of detail, when we are competent to do it.

Given this policy, it then becomes a matter of what behavior one is interested in modeling. I have chosen to sketch in a model of human behavior at a quite macroscopic level. In those instances where a working model has resulted, I'm quite confident that when the true neural circuitry is finally found, it will be found to carry out the same functions that are in the abstract model, with of course the realization that any given function could be carried out in innumerable different ways.

Considering the plasticity of any advanced reorganizable nervous system, it may be that the block-diagram level is the most detailed level at which any general model can be constructed. Individual organisms may accomplish the same functions in many different detailed ways -- the circuit you trace in one individual may not physically resemble the circuit you trace in a different one. If this is true, and I think it probably is, then the hope of building a model of "the" nervous system is in vain, because only at the functional level is one individual constructed like another. Even in the cockroach, what is found in the nervous system of one may differ in details of connection and neuron properties from what would be found from the equally-exhaustive investigation of a second cockroach. Yet both would behave the same way.

This is not to say that circuit-tracing is useless. When we understand one circuit's operation, and then a second one, and then another, we can come to see the general functions that are accomplished despite the variations in implementation. That would be one useful way to discover general functions like perception, comparison, and effector output.

But there is another way, and that is to start with behavior and ask what functions must be performed, whatever the implementation, in order that behavior occur as we observe it to occur. That's a different way of approaching the same functions, heading toward the same level of description at which there is some generality across individuals.

This is what I have tried to do. From a strictly external viewpoint, there is an infinity of different neural organizations that could create ahierarchy of control behaviors. But I have always tried to constrain the possibilities so they were consistent with what I know about the nervous system -- admittedly not much, but enough to eliminate many possibilities. The general concept of levels of perception is consistent with gross observations of brain functions at various levels. The general concept of feedback loops at each level is consistent with gross anatomical features such as recurrent collaterals. The process of comparison is consistent with the widespread occurance of places where these collaterals join downgoing "command" signals in a single location, with opposite signs. And so on. I have not proposed this hierarchical structure out of a vacuum, or simply by dreaming up organizations that might work. I haven't used models of individual neurons, although in BCP, 19 years ago, I indicated some building blocks for an analog neural model (long before the "neural network" movement). I have used information at the level where it appeared to bear on major functional questions.

Speaking of functional evidence, nobody else has mentioned this article in Science, so I will.

Van Essen, D. C., Anderson, C. H., and Felleman, D. J.; Information processing in the primate visual system: an integrated systems perspective. Science,_255_,419-423, 24 Jan 1992.

"To date 305 pathways [in macacque monkey's visual system] interconnecting the 32 cortical areas have been identified with modern pathway-tracing techniques. This constitutes nearly one-third of the number there would be if the network were fully interconnected. Hierarchical relations between areas have been assesed by the use of information about the cortical layers in which pathways originate and terminate. For some pathways the laminar pattern suggest ascending (forward) information flow from a lower to a higher area. These are generally paired with reciprocal pathways that have patterns suggesting feedback [note inversion (WTP)] from a higher to a lower area. Other pathways have patterns suggesting lateral connections between areas at the same level. Systematic application of these criteria leads to a hierarchy containing ten levels of cortical visual processing plus several additional stages of subcortical processing. The visual hierarchy is extensively linked to centers associated with motor control, other sensory modalities, and cognitive processing..." (p. 419).

So it seems I may have missed a couple of levels. On the other hand, this data is for macacque monkeys, and we must remember that these monkeys do not indulge in air pollution, destruction of environments and other species, genocide, and forms of behavior that suggest a limited mentality.

Avery Andrews' post deserves more attention than I can give it just now.

Best to all,

Bill P.

Date: Mon Feb 17, 1992 11:59 am PST
Subject: Individual Differences

[from Gary Cziko 920217.1330]

Bill Powers (920217.1000) says:

>Considering the plasticity of any advanced reorganizable nervous system, it
>may be that the block-diagram level is the most detailed level at which any
>general model can be constructed. Individual organisms may accomplish the
>same functions in many different detailed ways -- the circuit you trace in
>one individual may not physically resemble the circuit you trace in a
>different one. If this is true, and I think it probably is, then the hope
>of building a model of "the" nervous system is in vain, because only at the
>functional level is one individual constructed like another. Even in the
>cockroach, what is found in the nervous system of one may differ in details
>of connection and neuron properties from what would be found from the
>equally-exhaustive investigation of a second cockroach. Yet both would
>behave the same way.

Figure 10.2 on page 180 of Edelman (1988) provides a nice collection of illustrations of variability or parts of the nervous systems of various organisms. Of particular interest is "a schematic representatin of the branching sequence of fiber from an ommatidial receptor neuron of _Dapnia magna_ [I believe this is the "water flea"]. Neurons on the left . . . and right . . . of four genetically identical specimens . . . are shown" and there is a quite an amazing degree of variability in these four specimens. I assume that nonetheless the four genetically identical _Daphnia_ demonstrated more or less the same types of behavior.

Reference:

Edelman, Gerald M. (1988). _Topobiology_. New York: Basic Books.

Gary A. Cziko   Telephone: (217) 333-4382

Date: Mon Feb 17, 1992 12:50 pm PST
Subject: Branching Out

[from Gary A. Cziko 920211.1440]

For the last few months, membership ("subscriptions") to CSGnet has remained fairly stable at between 110 and 120. The last large increase resulted in an announcement describing CSGnet posted on the CYBSYS-L list. I am asking now if anyone out there knows of other lists or newsgroups which would likely include people who would find CSGnet of interest.

Part of the motivation for increasing "subscribers" is that at the last CSG meeting in Durango it was decided that we try to become an official Usenet Newsgroup. This would greatly increase our accessibility worldwide. But to do that, we need to receive at least 100 more "yes" votes than "no" in an e-mail poll, something which seems hardly likely given that we have only 110 to 120 current subscribers.

So let me know if anyone out there has suggestions. These should probably be sent to me directly unless they call for group discussion.--Gary

---

Gary A. Cziko   Telephone: (217) 333-4382

Date: Mon Feb 17, 1992 1:58 pm PST
Subject: Re: Branching Out

[Martin Taylor 920217 16:30]
(Gary Cziko 920211.1440] Shouldn't this be 920217?

I wasn't at Durango, so I can't comment on the argument as to making CSG-L into a usenet group. But I can say that one of the things that attracts me to CSG-L in contrast to any unmoderated usenet newsgroup is the seriousness and lack of flamage in the postings. Playfulness and criticism are fine. Usenet groups seem to have a tendency to get swamped by extreme versions of them, followed up by tirades asking the perpetrators to shut up.

Maybe it is elitist, but in my book that's not a bad word. I would prefer to keep CSG-L for people who want to learn about and to advance the topic rather than open it up to the brainless drop-ins on the tens of thousands of machines connected to the usenet.

Would Bill Powers' words of wisdom be heard over the noise, on usenet? I have my doubts.

Martin

Date: Mon Feb 17, 1992 2:05 pm PST
Subject: Re: Branching Out

[from Gary Cziko 920217.1600]

Martin Taylor (920217 16:30) says:

>(Gary Cziko 920211.1440] Shouldn't this be 920217?

Yup.

>Maybe it is elitist, but in my book that's not a bad word. I would prefer
>to keep CSG-L for people who want to learn about and to advance the topic
>rather than open it up to the brainless drop-ins on the tens of thousands
>of machines connected to the usenet.

Perhaps we should discuss the Usenet connection again before proceeding with this. I don't think many of the people at Durango had any experience with Usenet and so perhaps were not sufficiently wary of the "brainless drop-ins" on Usenet.

But in the meantime, I would still like to find other groups that should know we exist as a LISTSERV list. We can then decide on the Usenet connection. And if we do go Usenet, I believe we can always undo the connection between it and CSGnet whenever I can be convinced that it should be done.--Gary

Date: Mon Feb 17, 1992 2:50 pm PST
Subject: usenet?

For what it's worth, I agree with Martin Taylor's judgements on usenet. The Linguist mailing list has declined to become a news group for exactly such reasons.

Avery.Andrews@anu.edu.au

Date: Mon Feb 17, 1992 2:57 pm PST
Subject: Re: Branching Out

One possibility is to post announcements of the existence of CSGNet on groups such as comp.ai. Interested people can then sign on.

Avery.Andrews@anu.edu.au

Date: Mon Feb 17, 1992 3:29 pm PST
Subject: Andrews program

[From Bill Powers (920217.1700)]

Avery Andrews (920216) --

Your brief comments about your program are suggestive of more questions -- I hope you don't mind answering them patiently for this nonlinguist.

```
>The f-structures of `dog' and `cat' will appear in the whole f-structure
>as values of SUBJ and OBJ, and the tense value of the verb as value of
>TENSE, so the whole f-structure will be:
>
>  PRED  `Bite(SUBJ,OBJ)'
>  SUBJ   [PRED  `Dog']
>  OBJ    [PRED  `Cat']
>  TENSE  PAST
```

How does the program know that "bite" takes arguments of classes SUBJ and OBJ, whereas "dog" and "cat" don't?

I have similar questions about

```
>The rules say things like:
>
>   S  -> NP:SUBJ   VP.
>   VP    V  (NP:OBJ)  (>NP:OBJ).
```

Presumably, the rules have to be told whether something is an NP or a VP and the like before they can be applied. How are these identifications made?

I would love to get your programs, but I don't know what compression scheme you're talking about. Is this something available for PCs? Would any old VAX cluster or VMS system know about it? Help, please.

Best,

Bill P.

Date: Mon Feb 17, 1992 3:56 pm PST
Subject: LFG, compres

Re Bill Powers (920217.1700)

>How does the program know that "bite" takes arguments of classes SUBJ and
>OBJ, whereas "dog" and "cat" don't?

This is specified in the lexicon, which says that `bite's;'
contribution to the f-structure of the sentence it appears in
is

  PRED `Bite(SUBJ,OBJ)'

>Presumably, the rules have to be told whether something is an NP or a VP
>and the like before they can be applied. How are these identifications
>made?

More rules of the same type, e.g.

  NP -> Det N

bottoming out in the lexicon, which says what phrase-type each individual
word belongs to (think of this info as hanging off the ends of the
branches of the word-recognition tree).

`compress' is a unix utility - there might be VAX or PC programs
implementing compress/uncompress, but if there aren't I could
just mail the disks. But like I said, I'm not sure how much
people would get out if it without also having a suitable textbook
(like the one I'm working on).

  Avery.Andrews@anu.edu.au

Date: Mon Feb 17, 1992 4:28 pm PST
Subject: LFG files - confession of stupidity

I suddenly realize that what I should have done is compress the
LFG stuff with a PC utiity, and then stick it on the ftp server --
this oversight will be attended to shortly.

  mea culpa
  Avery.Andrews@anu.edu.au

Date: Mon Feb 17, 1992 6:42 pm PST
Subject: dotted line

[From Wayne Hershberger]

>(Mark Olson 920128)
>At first I thought I was biased against nonliving control systems.
>But now it seems that I find the problem with that nasty

>environment/organism dotted line. The line seems a little more
>"real" for the living than the nonliving.   When the temperature
>drop of the room causes (ontologically) the furnace to go on, I can
>explain it with physics and the line itself becomes a designation
>without real ontological status. That line seems "more real" with a
>living system.

>(Wayne Hershbeger 950205)
>To me, the organism-environment interface is every bit as artificial
>(nasty) as the mechanism-environment interface.

(Mark Olson 920216)
>So are you saying that the line should not be there. That the
>organism/environ interface is really no different than the mech/environ
>one? Causality will creep in really easily if I accept that. Are events
>"out there" partial causes of our behavior in the same way the the
>temperature of the room is a partial cause of the furnace going on? I
>still can't reach a satisfying answer on this. ANy suggestions?

The dotted line serves a purpose, but it does not represent the control system's outer bound. Rather, it locates the transducers we call receptors and effectors that JOIN the ecological system's two major subdivisions. As for partial causes, I recommend you call them what they are (influence) and the problem should become manageable. Influencing, determining, and controlling are distinctly different concepts/phenomena. Bill Powers posted an excellent essay last year on this topic; did you see it?

I think you are wrestling with a very important issue, but I am also confident that you will get it sorted out to your satisfaction. Further, I think that my chapter in _Volitional_action:_Conation_and_control_ might prove helpful to you. What is your US mailing address? I'll send you a copy.

Consider the canonical control loop, drawn below, with its two inputs (inputs to the loop): one is the reference signal, the other is the disturbance.   Note that the output of the organism is a joint effect of C and R. Note also that the organism's controlled input is a joint effect of D and O. Although both D and O jointly determine C, their individual influence on C is effectively nil because D elicits compensatory output whose influence on C is equal and opposite to D's influence on C. So, although the value of C is a joint effect of D and O, the value of C can not be altered by changing D; C can be changed only by altering the value of R--providing that the control loop keeps error (R-C) effectively nil.

```
                              Reference input
                                    |
                                    |
                        ____________|___________
Controlled input, C ->  |         (R-C)         |
                        |                       |
                        |                       |
organism                |                       |
- - - - - - - - - - - - | - - - - - - - - - - - |- - - - - - - - - -
environment             |                       |
                        |                       |  output, O
                        |         (O+D)         |
```

```
                        |_______________________|
                                   |
                                   |
                                   |

                              Disturbance
```

Warm regards, Wayne

Date: Mon Feb 17, 1992 10:10 pm PST
Subject: Re: branching out; BEERBUG

[From Tom Bourbon 920217 23:50]

BRANCHING OUT. Martin Taylor and Avery Andrews, both [920217]:
Second!
BEERBUG: Greg Williams [920217] asked why the debate on Beer's bug is "threaetning" to PCT people. I didn't realize my three-part post contained a hint that I was threatened. I am not. I always cast a keen eye toward material that purports to reveal the physiological bases of anything to do with perception or behavior. Perhaps it is reaction formation to my student days when I almost literally had to "pledge allegiance to physiology" as the key to understanding any and all topics in psychology. The *wish* that we might have such explanations is ancient; the *hope* that data sufficient to the task will soon be in hand seems eternal. Beer (admirably) reports that there are few data to support many features of his model, and for all of the subsystems that regulate the important behaviors he labels "exploration," there are NO DATA, or so he reports.

In the face of Beer's repeated candid reports, I am somewhat puzzled at your frequent claims that he works closer to some important level of detail than do PCT modelers. I don't think my puzzlement is an indication that I feel threatened. Frankly, during the long period when I was confined to watching the debate unfold, I suspected that you might be playing devil's advocate -- somewhat like your brilliant performance at one of the "mini-conferences" at Bill Powers' place in Illinois, where you assumed the role of B. F. Skinner. This time, I thought you might be playing a wolf in bug's clothing. On the evidence cited by Beer, how do you conclude that his "exploration" components, for example, are modeled close to physiological data? Enlighten me.

Within a few months, I will be moving to Galveston, Texas, to join Andy Papanicolaou at the magnetoencephalography lab there. We intend to do as much work on and with PCT as we can. One of our highest priorities is to recruit single-cell physiologists from among the people in the Division of Neurosurgery at the medical school, and from the Marine Biological Institute, just across the street. We hope to convince some of them that it might be worth their while to explore some "simple" nervous systems from the perspective that the creatures control some important variables in themselves and their environments. We believe such an exploration might yield interesting results. Stay tuned.
INDIVIDUAL DIFFERENCES. Gary Cziko [920217] The citation on differences in nervous systems of conspecifics looks good.
Thanks for sharing the information

Tom Bourbon   <TBourbon@SFAustin.BitNet>

Dept. of Psychology
Stephen F. Austin State Univ.
Nacogdoches, TX 75962   Ph. (409)568-4402

Date: Tue Feb 18, 1992 7:07 am PST
Subject: Re: WTP & TB on Beer's Bug

From Pat & Greg Williams (920218)

>Bill Powers (920217.1000)

>I would use neural data in my models if I could. I'm prevented in part by my
>lack of access to literature (in any easy way)...

Of course this is not a good excuse -- to play the game, one must play it.

>... and my priorities.

This IS a good excuse. Let's examine your priorities.

>I'm also prevented by many of the factors Tom mentions: lack of data that's
>better than an educated guess.

With respect to detailed models of high-level human control structures, we agree that data are lacking. We think considerably more data are available with respect to detailed models of various "simple" organisms.

>The data that are available show many of the biases of neurological
>researchers -- the emphasis on thresholds, for example, and the use of on-off
>sensors and "state" output variables.

This sounds like a cop-out. It smacks of "you won't play by MY rules, so I'll take MY ball and go home." We don't believe that ALL (or even MOST) of the data available from neuroethological investigations are so skewed as to be unusable by PCTers, unless those PCTers willfully reject it because it doesn't fit their preconceived notions.

>Those are purely theoretical ideas, and they carry the flavor of belief that
>single impulses are significant and represent digital variables -- even though
>"neural currents" are used. The concept of continuous neural functions has
>been adapted to allow the use of digital variables after all!

This is a good example of a willful preconception. You appear to distrust data showing "digital variables" because you have a bias against non-continuous control processes at the philosophical level (because you want to contradict notions of discrete stimulus-and-then-response). CONTROL DOESN'T HAVE TO BE CONTINUOUS.

>There's some confusion in Beer's model, I think, between thresholds of
>firing and thresholds of response in terms of frequency inputs and outputs.
>Even if there are thresholds for generation of a single spike, there need
>be no thresholds when you measure output frequency as a function of input
>frequency (although there will be nonlinearity at the lowest input rates,
>and perhaps some threshold effect).

Please explain what you mean here, with emphasis on: so what?

>Looking at the plots in your program, I see that the current outputs tend to
>jump from one value to another, although the voltage measures are smooth. I
>presume this is the threshold effect showing up, converting an analog
>relationship to a digital one by using high gain and a threshold -- much as
>TTL logic works. But this is a matter of how you adjust the parameters of the
>neuron, which I presume is not an empirically-based choice. Different choices
>would have led to an analog-computer model.

What is going on is physical causality. The membrane of a modeled neural cell has a (lumped) resistance and capacitance across it. In such a circuit, the current through the membrane can change instantaneously (in this case, via "synaptic" inputs which are modeled as current changes), but voltage (of which the cell's firing frequency is a thresholded, piecewise-linear, and saturating function) CANNOT change instantaneously -- it can only change via first-order dynamics. We think this is a quite realistic model for neuronal function. It is basically analog; the essentially on-off firing patterns of some cells reflect their input patterns, not thresholding. For example, the leg angle sensors are modeled to have 10 nA of current input if the angle is greater than a certain value, 0 nA otherwise; to make the sensors analog, you could simply make I(in) a continuous function of leg angle. But would you want to do that? Better look at the available data first, and not just do what would make your preferences easy to fulfill. (One place where such on-off firing is extremely well documented is in the lobster stomatogastric neural "system" -- various neurons tend to either fire at an essentially constant rate or not fire at all, coordinative with the churning of the gastric "mill.") An aside: where we think Beer's neuron model is LEAST realistic is in his use of the membrane time constant (making it long) to extend the activity of a particular cell over time, whereas this probably takes place in real cockroaches via recurrent circuits of more than one neuron. And, of course, Beer's individual neurons each "stand for" the activity of MANY neurons in the real cockroach.

>I'm not totally convinced that the "oscillator" concept is the only good one,
>however. Applied to human walking, it doesn't explain how walking can slow
>down, stop, reverse, go forward again, etc.

Beer's gait generator can be slowed, stopped, started, and speeded by altering the firing level in one neuron (LCS). Beer had reversing in an early model, but took it out of the final one. Maybe it needs to be made explicit that when Beer's bug is "hungry" and gets to "food," it slows and sometimes stops while "eating," then speeds up, leaving the "food" patch.

>It wouldn't do for a model of the way Michael Jordan moon-walks!

It certainly wouldn't. Much room for improvement, we suppose!

>My model is empirically based, perhaps considerably more so than Beer's.

That is too big a "perhaps" for us. Here is where we have a fundamental disagreement with you, which ultimately is an opinion. Putting our view as compactly as we can, we claim that the testability of detailed conjectures in your model doesn't come close to the testability of those in Beer's model. This doesn't mean that there aren't some testable-to-some-extent features in your model, and it doesn't mean that all features of Beer's model are testable. But there is, overall, a great contrast in testability -- a contrast which we think will actually increase over the next few years as the ratio of data-which-could-be-used-to-test-models for control structures of "simple" organisms to that for control structures of "complex" organisms continues to grow. Testability is not everything. But we think it weighs heavily in the

choice by scientists about what they should be attending to -- and especially
in what should be published and disseminated. Note once again that we are NOT
claiming that your model is wrong and that Beer's is correct. Unfortunately,
we get the impression that your own commitment to the (untested) "probable
correctness" of your model is leading you to ignore contrary possibilities.
Beer explicitly states that his model is a "first shot" and he appears to be
attempting to improve it, based on increasingly available data.

>"Empirical" doesn't necessarily imply dissection; it means only building
>explanations on observations. I chose to start with observed behavior, and
>to try to characterize organizations that will produce what we observe. Themain
>phenomenon we observe is control. As I showed in my last post, it's
>possible to deduce the presence of a control system strictly on the basis
>of externally-measurable variables, without a block diagram or even a
>nervous system.

In other words, your GLOBAL (we've also used the word "lumped") control model
IS testable. We agree. Furthermore, we think that global control model
"passes" the tests (with the usual caveats about scientific testability).
That's why we said sometime back that the lumped model won't be displaced
(although it might be added to, say with ideas from adaptive control).

>In a similar way it would be possible to show that a hierarchy of control
>exists, by observing cases in which a variable already characterized as a
>control phenomenon becomes the means of controlling another variable that is
>in part a function of the first one.

As the details are brought in, the case starts becoming much less cut-and-
dried, in our opinion. We think Beer's bug instantiates your lumped model (and
the world needs to be made aware of this!), but uses a fundamentally different
detailed -- even at the level of detail of functional block diagrams -- method
of hierarchical (AND sometimes not strictly hierarchical) control. Knee-jerk
objections to the fundamental differences are what we dislike.

>I have chosen to propose a simple functional model of the insides of the
>behaving system, based partly on my experience with electronic control
>systems, that would produce the kinds of control behavior we can see. I've
>said many times that this is a "canonical" model, meaning that it is a form
>that's equivalent in function to many other possible forms.

It isn't equivalent to ALL other possible forms. Our deepest misgivings arise
when PCTers choose detailed models ONLY in conformity to YOUR "canonical"
model. There are no good reasons, we think, for discounting alternatives to
your "canonical" model. And there are some good reasons for considering
alternatives, not the least of which involves making PCT seem less dogmatic.

>I have, as you can see, used such neurological data as I have found, although
>it would be futile to try to make the details any finer, given the state of
>neurological knowledge.

Agreed. There just isn't much data relevant to your detailed model. That's
half of our argument. The other half is that there is considerably more data
relevant to models of various "simple" organisms. Again, this is ultimately a
matter of opinion. If we enumerate such data, some will say it is tainted by
S-R thinking, it is "model" rather than "fact," etc.

>And there is no assurance that the connections shown in the model are those
>that actually exist, according to Tom's report.

That depends on what you mean by "assurance" -- we think there IS some
assurance -- considerably more assurance than that your detailed model's
"higher-level error alters lower-level reference signal" scheme actually
exists above, say, 4th order in humans.

>Another difference from my model -- an important structural difference from
>the canonical model -- is the cross-connection between systems at the same
>level. I've mentioned that possibility many times, but have never seen a
>good way to incorporate it into the model. Maybe Beer (or those whose work
>he used) has shown a good way. I will certainly try it out.

Great! An example of learning from other models....

>The implementation of such models using pseudo-neurons, however, is basically
>a trick, because only the mathematical relationships represented by those
>neurons actually matter -- until we can say that THIS is the true circuit
>diagram, THESE are the true parameters of the individual neurons, and THOSE
>are the correct models of the sensors, muscles, and physical consequences of
>action.

Basing the equations on data about how real neurons work is indeed some trick
-- the data provides CONSTRAINTS on those equations which would otherwise be
much easier to, armchair-fashion, devise. One can go from the neuron data to
those equations easily enough, but one cannot go from unconstrained equations
to the underlying physiology (and thus some onlookers will worry about
"miracles" having been postulated -- that WOULD be prestidigitation).

>... we can't know which of the infinity of possible designs is actually used
>until we no longer have a choice about how to interconnect the neurons and how
>to adjust their parameters.

We don't think one needs to know EVERYTHING about the functional organization
of a system to begin modeling. One can eliminate large classes of models and
settle in on a few classes.

>While we do still have a choice, it's much more realistic to just draw
>boxes and assign functions to them, looking for a design that will
>economically accomplish the behavior we observe from outside. How the
>function in each box is accomplished in neurons is a matter for
>investigation at another level of detail, when we are competent to do it.

We are reminded of the drunk looking for his lost car keys under a
streetlight, even though that isn't near where he lost them, because "the
light's better here." Nevertheless, your point is quite reasonable IF you
remain open to the possibility of alternative designs all capable of
"economically accomplishing the behavior."

>In those instances where a working model has resulted, I'm quite confident
>that when the true neural circuitry is finally found, it will be found to
>carry out the same functions that are in the abstract model, with of course
>the realization that any given function could be carried out in innumerable
>different ways.

Our confidence falls as you get to more detailed modeling.

>Considering the plasticity of any advanced reorganizable nervous system, it
>may be that the block-diagram level is the most detailed level at which any
>general model can be constructed. Individual organisms may accomplish the
>same functions in many different detailed ways -- the circuit you trace in

>one individual may not physically resemble the circuit you trace in a
>different one. If this is true, and I think it probably is, then the hope
>of building a model of "the" nervous system is in vain, because only at the
>functional level is one individual constructed like another. Even in the
>cockroach, what is found in the nervous system of one may differ in details
>of connection and neuron properties from what would be found from the
>equally-exhaustive investigation of a second cockroach. Yet both would
>behave the same way.

Making models at the neuronal level would be a good way to investigate and understand the limits on such variation.

>From a strictly external viewpoint, there is an infinity of different neural
>organizations that could create a hierarchy of control behaviors. But I have
>always tried to constrain the possibilities so they were consistent with what
>I know about the nervous system -- admittedly not much, but enough to
>eliminate many possibilities.

In sum, we think that many more possibilities could be eliminated by making models of control structures of "simple" organisms which are subject to both behavioral and physiological empirical constraints.

>I have used information at the level where it appeared to bear on major
>functional questions.

We think it is appropriate to move beyond that level now.

>Tom Bourbon

>Greg Williams [920217] asked why the debate on Beer's bug is "threaetning" to
>PCT people. I didn't realize my three-part post contained a hint that I was
>threatened. I am not.

From Greg: I didn't say that I thought YOU were threatened. You said the following in one part of your post: "I do not say this to be critical of Beer, but anyone who feels threatened due to a belief that Beer argues from hard data in physiology, exclusively, need have no such fears. Data of that kind are not available." I took this as your feeling that SOMEBODY was threatened. And I asked WHO? As I said, I, as a PCTer, don't feel threatened.

>On the evidence cited by Beer, how do you conclude that his "exploration"
>components, for example, are modeled close to physiological data? Enlighten
>me.

From Greg: I don't conclude that. The important "closeness" to physiological data for Beer's bug is with respect to the individual neuron models and the gait control circuitry. Nevertheless, I think it is possible to get "closer" to physiological data in modeling exploratory behavior of certain "simple" animals -- much "closer" than is possible in modeling high-level control circuitry in humans.

>One of our highest priorities is to recruit single-cell physiologists from
>among the people in the Division of Neurosurgery at the medical school, and
>from the Marine Biological Institute, just across the street. We hope to
>convince some of them that it might be worth their while to explore some
>"simple" nervous systems from the perspective that the creatures control some
>important variables in themselves and their environments. We believe such an
>exploration might yield interesting results. Stay tuned.

HOORAY!!!

Pat & Greg

Date: Tue Feb 18, 1992 8:26 am PST
Subject: Re: LFG, compres

[Martin Taylor 920218 10:30]
(Avery Andrews 920218 10:46 the wonders of modern science!)

>
>`compress' is a unix utility - there might be VAX or PC programs
>implementing compress/uncompress,

There is a public domain (or maybe it is shareware) version for the Mac
called MacCompress. Check on one of the archive servers for it if you
use a Mac.

Martin

Date: Tue Feb 18, 1992 9:49 am PST
Subject: SOME INFORMATION PLEASE

<<<<<<<<<<<<<<<<<FROM CHUCK TUCKER 920218>>>>>>>>>>>>>>>>>

I would really appreciate two pieces of information:

1. The exact and complete citation for Randy Beer's book

2. Information on how to obtain the program for Beer's Bug

Thanks in advance.

Chuck

Date: Tue Feb 18, 1992 10:30 am PST
Subject: Re: Williams on Bourbon and Powers on Beer

[From Rick Marken (920218)]

First, let me say that I agree with Martin Taylor and Avery Andrews (and
Tom Bourbon too) about the potential downside of making CSGnet a newsgroup.
I think that it probably would be best to just announce the availability of
CSGnet to relevent newgroups -- sci.psychology, comp.ai, robotics, etc.

Greg Williams (920217), in response to Tom Bourbon, says:

> Yes, the facts and conjectures
>upon which the bug model is based might be skewed by S-R thinking. Still,
>Bill's "functional"-level modeling of humans is a good deal more distant from
>the data than is Beer's "neural"-level modeling of bugs. ALL models are
>hypothetical, but some are more so than others.

So what your saying is that Beer is using realistic parts to build a model of
a non-existent phenomenon (S-R behavior)? My whole problem with Beer (and

Bill and Tom have tried to make this point as well) is that he has not identified the variables the bug controls. He treats behavior as a generated output (regardless of how it is generated). I don't see the value of having detailed, realistic models of a phenomenon that does not exist. It strikes me as having the same value as taking some well known astronomical components (comets, orbital eccentricities, whatever) and piecing them together to explain the Biblical account of Genesis (boy, am I going to get in trouble with THAT one -- loose canon and now INFIDEL).

> With regard to the PCT
>"agenda," as Mary calls it, I think the fact that you and I disagree about the
>relative physiological grounding of Beer's and Bill's models isn't as
>important as the apparent fact that modeling ala Beer is generating a lot more
>interest than modeling ala Bill. I suggest that the implicit difference in
>both actual and potential empirical support for the two endeavors is part (but
>certainly not all) of the reason.

I think the reason for the greater interest in Beer than Bill is the the reason for the greater interest in fuzzy logic, expert systems, non-linear systems, etc than PCT -- it's because most students of living systems assume that organisms generate output rather than control input. They assume that behavior is the kind of phenomenon that is consistent with their models and methodologies. This is why there are probably no more than 5 psychologists in the country who work in the context of PCT. If knowledge were a popularity contest, I'd jump on the Beer bandwagon immediately.

Tom says:

>>Beer picks target "behaviors" and designs plausible component circuits to
>>produce exactly those behaviors -- one at a time --then he assembles the
>>components from his software parts kit into a complex system.

Greg asks:

>How does this differ from (say, Bill's arm-) modeling in PCT?

What Tom meant is that PCT tries explicitly to model CONTROL. Beer's models do not. This is THE big difference to me. The fact that some control emerges out of the Beer model is nice; but if that was not the goal of modelling then I don't see how you can evaluate the model (except by continuing to get surprise results). Also, if the "neuro-physiologically correct" components do not produce a system that controls the variables that it is known to control then who is right? If the bug controls the angular position of its leg (for example) and we find no circuits that allow this kind of control then do we say that the bug isn't really controlling that variable? It seems to be that, by going with the "neuro-physiology" we could end up like the (possibly apocryphal) engineers who found that it was impossible for bumble-bees to fly.

>My overall point is that there are several organisms with CNSs to be modeled.
>At this time, I think it makes good tactical sense to model the simpler ones.

That's fine. But I would like to know more about the control abilities of the simple organism's whose CNSs are so well known (if they really are). Maybe this data is available but I saw nothing in Beer's book that suggested that he was trying to design into P. computatrix the ability to control any variables.

> Again I say why must PCTers cling to ONLY Bill's (detailed)

>model, given the paucity of data (but some negative data)?

It's not the model I cling to -- it's the phenomenon. The fact that we can model the phenomenon gives confidence that it can be produced by a real system -- this takes the "metaphysical mystique" from the phenomenon of purpose (control) which many scientists once rejected as fancy. Bill and I have built models where higher order systems alter properties of lower order output functions rather than their references (this was not part of the original model) because it looked like this is what happens. I am happy to use Beer-type components to model control. But that's what I would want to use them for -- to produce control, not just some "output" (like a "gait") that happens to catch one's fancy. First, I want to know what variables are involved in "gait" and whether they are controlled. Then I would start to model.

Pat & Greg Williams (920218) in reply to Bill Powers say:

> we claim that the testability of detailed conjectures in
>your model doesn't come close to the testability of those in Beer's model.

What would constitute a test of Beer's model? If the tests are at the physiological level, then you are probably correct. But aren't there to be behavioral tests? Is Beer's model successful if the components behave correctly even if the simulated bug can't balance on uneven terrain?

> Note once again that we are NOT
>claiming that your model is wrong and that Beer's is correct.

I think they are orthogonal. Beer is not TRYING to model control. The model controls -- but this seems to be a surprising (and satisfying) side effect.

> Unfortunately,
>we get the impression that your own commitment to the (untested) "probable
>correctness" of your model is leading you to ignore contrary possibilities.

The "contrary possibilities" seem to be in the building blocks. If one of the "contrary possibilities" is that organisms generate outputs, then we know already that they do -- these are accidental side-effects of control. I'm just not trying to model those side effects -- I want to model the phenomenon of control (although doing so might involve including some uncontrolled components).

>In sum, we think that many more possibilities could be eliminated by making
>models of control structures of "simple" organisms which are subject to both
>behavioral and physiological empirical constraints.

I AGREE!! But don't forget those BEHAVIORAL CONSTRAINTS. That is what I think is seriously missing from the Beer work. He may think he is working under behavioral constraints, but his concept of behavior is very different than mine. That makes it impossible to evaluate the model (from my perspective) other than as an interesting attempt to engineer a mobile robot made out of components that obey fairly strict physiological constraints.

Again, I am not interested in bashing Beer. He has done lots of very interesting work. I just don't immediately see its relevance to my interest in purposive behavior. But maybe you (or he) could help clear that up for me.

Best regards

Rick

Date: Tue Feb 18, 1992 1:50 pm PST
Subject: Beer b/b info

From Pat & Greg Williams (920218-2)

Randall D. Beer, INTELLIGENCE AS ADAPTIVE BEHAVIOR: AN EXPERIMENT IN COMPUTATIONAL NEUROETHOLOGY, Academic Press, San Diego et al., 1990, ISBN 0-12-084730-2.

Nervous System Construction Kit Version 3.0 (for IBM computers with EGA or VGA video) is available for $10.00 cash/check/m.o. prepaid from: Pat & Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328. Please specify 360KB, 1.2MB, or 720KB disk format. Turbo C source code is included. May be freely distributed and modified.

Pat & Greg

Date: Tue Feb 18, 1992 3:27 pm PST
Subject: Lagnuage; BEERBUG

[From Bill Powers (920218.0900)]

Avery Andrews (920217) --

For me, a "ZIP" file or any self-extracting file would be best. LHARC is another.

>>How does the program know that "bite" takes arguments of classes SUBJ and
>>OBJ, whereas "dog" and "cat" don't?
>
>This is specified in the lexicon ...

OK, this takes us another step. What are the processes that build the lexicon? If you do it yourself, how do you know that "bite" is of a class that takes a SUBJ and OBJ class of arguments?

Similarly for

> NP -> Det N
>
>bottoming out in the lexicon, which says what phrase-type each individual
>word belongs to ...

Same question: what must be known (perceived) in order to know what phrase-type is appropriate?

I'm asking these questions seriously. They probably call for some introspection, perhaps informed by HCT.

(Question all authors hate) When do you think this textbook of yours might
go to press?

Greg & Pat William s (920218) --

Looks like we're in a fight over whose model (or modeling philosophy) is
best. At least let's keep it clean:

me:
>>I would use neural data in my models if I could. I'm prevented in part by
>>my lack of access to literature (in any easy way)...

you:
>Of course this is not a good excuse -- to play the game, one must play it.

The only appropriate response to that, were I to accept the spirit of the
comment, would be to say "Gee, I'm sorry, I guess I can't play with the big
guys then, so I apologize and will shut up." Is that more or less what you
had in mind?

me:
>>... and my priorities.

you:
>This IS a good excuse. Let's examine your priorities.
Thank you for your evaluation. But what if I flunk your exam? I think I'd
better keep my priorities to myself. Presumably, I'm behaving according to
what they are.

>>The data that are available show many of the biases of neurological
>>researchers -- the emphasis on thresholds, for example, and the use of
>>on-off sensors and "state" output variables.

> This sounds like a cop-out. It smacks of "you won't play by MY rules, so
>I'll take MY ball and go home."

I suppose it does, although if the sensor responses and output variables in
question aren't actually binary in nature, my comment might sound less
petulant.

>We don't believe that ALL (or even MOST) of the data available from
>neuroethological investigations are so skewed as to be unusable by PCTers,
>unless those PCTers willfully reject it because it doesn't fit their
>preconceived notions.

Pretty strong accusations, here. How come you associate with people like
this? Are you saying I should play by HIS rules or HE'LL take his ball and
go home?

>>The concept of continuous neural functions has been adapted to allow the
>>use of digital variables after all!

>This is a good example of a willful preconception.

Yes, I thought it was. Oh, you meant ME!

> You appear to distrust data showing "digital variables" because you have
>a bias against non-continuous control processes at the philosophical level

I don't distrust DATA showing digital variables. I distrust the ASSUMPTION of digital variables as a SUBSTITUTE for data. Are there actually any digital variables in the (real) cockroach's machinery? Tell me there are, and I'll accept it.

>>There's some confusion in Beer's model, I think, between thresholds of
>>firing and thresholds of response in terms of frequency inputs and
>>outputs.

>Please explain what you mean here, with emphasis on: so what?

Output frequency is a function of the difference between a local chemical potential near the axon hillock and the potential at which a spike is initiated. If the local potential is maintained above the firing potential, the neuron will recover in a certain time and fire again: it will fire at a certain frequency. As the local potential varies, the recovery time varies and thus the firing frequency varies.

The local potential depends on the net difference between all excitatory effects of incoming signals and all inhibitory effects. The amount of inhibition determines the amount of excitation needed to produce net excitation, and thus a non-zero firing rate. So inhibitory inputs create the appearance of a threshold of frequency response which is not the same as the voltage threshold at the axon hillock. I would prefer to treat apparent thresholds as primarily an effect of inhibitory inputs.

A continuous digital signal is simply an analog signal that remains steady at a high frequency or at zero frequency, passing rapidly from one value to the other. Simply looking at such a signal doesn't tell you whether it has only two stable states, or represents an imbalance between relative large excitatory and inhibitory effects. If one looks only at effects of input signals that are relative large and change by large amounts, the imbalance will appear only as a high output frequency or a low (or zero) frequency. To see if a high-gain amplifier is involved, one would have to adjust the net excitation delicately over a small range, looking for a range over which the output frequency varied smoothly with the input frequency. Such a neuron could be part of a control system with a very high loop gain. With the loop closed, feedback would keep the neuron within its proportional-response range, and we would see continuous variations. With the loop open, we would see what looks like an on-off signal. That is the so what.

>What is going on is physical causality. The membrane of a modeled neural
>cell has a (lumped) resistance and capacitance across it. In such a
>circuit, the current through the membrane can change instantaneously (in
>this case, via "synaptic" inputs which are modeled as current changes),
>but voltage (of which the cell's firing frequency is a thresholded,
>piecewise-linear, and saturating function) CANNOT change instantaneously
>-- it can only change via first-order dynamics. We think this is a quite
>realistic model for neuronal function.

I agree; it's a good starting-place. A comment: With a steady input signal, voltage will rise until the voltage lost between input impulses matches the voltage gained from each impulse. This is the steady-state input gain. If capacitance is raised and leakage resistance is lowered, the time-constant will remain the same while the apparent input gain decreases. So by adjusting capacitance and resistance, it is possible to vary gain and time constant independently. The overall gain factor is thus redundant.

What we get from this model is a building block that can be adjusted so it produces an output frequency depending on an input frequency with any gain and time constant we please. I would prefer, by my argument above, to leave thresholding up to inhibitory inputs, but that's not critical.

Now, how shall we adjust these parameters and interconnect the neurons? One approach is the design approach: we ignore the data and simply produce a design that seems to behave correctly, in terms of our understanding of behavior. Another is to build the circuit strictly by reference to measurements of input-output functions of actual sensors, neurons, and effectors, using only the connections found actually to be there. I judge that Beer is somewhere between these approaches, closer to the first than to the second (after all, the entire pacing system is imaginary).

With the pure second approach, theory is unnecessary, and so are observations of behavior. If the model system has the right properties and the right input and output devices, and the right environment, it can't help behaving as the real system behaves, at least barring learning and higher-level effects. Of course you would have to do this for every individual organism, as no two of them are wired alike. The organism would not be in working order after you finished, so the result would be of only historical interest.

With any approach between pure design and pure circuit-tracing, the neural data can only suggest, not guide. It is equally important to know what behavioral phenomena the real system produces that the model also has to produce. It's extremely important to know how variations of output affect the inputs, because apparent causality, even inside the neural system, is greatly affected by feedback. One can be misled by an apparent relationship that simply disappears when the external feedback path is cut. Piecewise measurements treat inputs to components as independent variables; when a closed-loop system is operating, both input and output are dependent variables. So behavioral data and a realistic model of the external part of the loop are essential, the more so as we work farther from the limit of quantitative measurements and circuit tracing and more toward the design end.

If I were designing the edge-tracking part of the cockroach (as I will be doing pretty soon, I hope), I would make the antenna sensors proportional, so as the antenna's nominal position came within a certain range of the obstacle the signal would begin to grow (this is equivalent to making the antenna a little springy and having the sensor record force). As the signal grows, it would alter the direction of movement -- I don't know yet how the signals would affect the pacing circuit (which I will use) to bend the path. The result would be to keep the force on the antenna close to zero, and the path would follow the edge. I don't know what to do about inside corners -- that will probably take a high-level system or some "reorganizing" effect.

My cockroach will follow an edge smoothly, whereas yours keeps bouncing away and returning to the edge. As there appear to be no data concerning the actual response of the antenna touch sensors, our models would be on equal footing as far as the neural facts are concerned. We would then have to turn to the actual cockroach. How does it behave when following edges: like your model, like mine, like something in between, or like something else entirely?

>Beer's gait generator can be slowed, stopped, started, and speeded by
>altering the firing level in one neuron (LCS).

Wonderful! So we can treat the whole gait system as an output function, and hook it up to sensors having to do with forward progress. How does the model alter direction of walking?

>Putting our view as compactly as we can, we claim that the testability of
>detailed conjectures in your model doesn't come close to the testability
>of those in Beer's model.

Do you mean actually, or in principle? I thought Tom said that even the existence of the pacing circuitry in Beer's model was unsupported by observation. It's pretty hard to compare in-principle verifiability of two theories. What beside the properties of individual neurones (highly adjustable in Beer's model) and the existence of sensors and effectors of unmeasured properties, is available for verification of Beer's model?
Much of the verifiability of my model is behavioral. In tasks ranging from elementary to somewhat complex, my model generates predicted traces of limb movements that match actual limb movements with an error of five percent or better. How closely do the movements of the cockroach in Beer's model match the movements of real cockroaches?

In the arm model, the role played by stretch and tendon receptors is taken fairly directly from the literature, including tonic and phasic aspects of the stretch reflex, connection to and from the spinal motor neurons, and the signs of the effects. The muscle model, while simple, does incorporate the mass-spring properties correctly if linearly. The visual part of the model bypasses any detailed modeling and just uses the effects that are required: comparison of finger position and target position, however sensed, and conversion to kinesthetic reference signals, however accomplished. The external (visual) geometry is correct in most details. The overall behavior is reasonably good; the model can reach out and touch a target, and follow it around with or without gravity, although the behavior under extreme conditions is probably not correct (movements at maximum speed from one configuration to another, for instance). How does Beer's model fare under extreme conditions?

Why do you insist that my model is conjectural where Beer's is not? I have more supporting detail for the kinesthetic parts of my arm model than Beer has for his cockroach. The physical part of this model, thanks to you, is incomparably more realistic than Beer's treatment of the world in which his cockroach lives and the mechanics of the body. I could convert the equations for the control systems in the kinesthetic part of the arm model to Beer neurons, but what would be the point? They would be the same equations. I would select parameters for the neurons that would make them the same. This wouldn't prove that I had the circuit right, only that one circuit made of those components would implement the equations I've used.

>There are no good reasons, we think, for discounting alternatives to
>your "canonical" model. And there are some good reasons for considering
>alternatives, not the least of which involves making PCT seem less
>dogmatic.

I'm all for that. I will remain dogmatic about one thing: if control phenomena are involved, the model has to reproduce them accurately. If the real control is poor and the model controls poorly in the same way, bravo for the model. If real control is good and the model controls poorly, then I don't much care what's in the model: it's wrong. Go away, make it work right, and we'll have another look. I'm not going to change the CT model JUST to make it look less dogmatic.

>There just isn't much data relevant to your detailed model. That's half of
>our argument. The other half is that there is considerably more data
>relevant to models of various "simple" organisms.

And the other other half is that the behavioral data in any case has to match. As one of my highest priorities is to understand human behavior, I have only a general sort of interest in how other organisms, particularly very simple ones, work. I doubt that human beings contain a system for coordinating six legs, and I doubt that a cockroach would do very well at the Winter Olympics. What we learn from a cockroach's circuitry is not going to tell us much about anything but cockroaches, especially when we're given freedom to invent any properties for their neurons that we please, and create any physical properties of their bodies and environments that seem to work.

>>I have used information at the level where it appeared to bear on major
>>functional questions.

>We think it is appropriate to move beyond that level now.

I have a better idea. Why don't we abandon this stupid argument? You do what you're interested in, I'll do what I'm interested in, we'll collaborate where the domains intersect, what will come of it will come of it, and we'll still be friends. You got a higher priority?

Best to all

Bill P.

Date: Tue Feb 18, 1992 4:53 pm PST
Subject: Re: Rick on... bug

From Pat & Greg Williams (920218)

>Rick Marken (920218)

>So what your saying is that Beer is using realistic parts to build a model of
>a non-existent phenomenon (S-R behavior)?

No, he is building a model organism which CONTROLS (as in control theory).

>My whole problem with Beer (and Bill and Tom have tried to make this point
>as well) is that he has not identified the variables the bug controls.

We don't agree. We think that Beer is not completely clear about the difference between "actions" and "behaviors" as elucidated by Bill, but he does sometimes make the distinction, and identifies some controlled variables.

>He treats behavior as a generated output (regardless of how it is generated).

Again, there is not complete consistency with PCT, but we think Beer is close. You have Beer's book, but for those (including Bill Powers) who do not, consider the following quotes the book (even the Devil can quote scripture -- by the way, we loved your (unintentional?) pun: "loose canon and now infidel" -- and consider substituting PCT's "action" for Beer's "behavior" and PCT's "behavior" for Beer's "task." We think he doesn't come off as so divergent from PCT gospel.

xvi - "To me, this penchant for ADAPTIVE BEHAVIOR is the essence of intelligence: the ability of an autonomous agent to flexibly adjust its behavioral repertoire to the moment-to-moment contingencies which arise in its interaction with its environment."

xvi - "The design of this artificial insect is based in part upon specific behaviors and neural circuits drawn from several natural animals. Its behavioral repertoire includes locomotion, wandering, edge-following, and feeding."

12 - "What I am trying to emphasize is the way in which the behavior of an intelligent agent engaged in ongoing interaction with its environment is continuously adjusted to the changing internal and external circumstances of that interaction in such a way as to achieve the agent's objectives. Strictly speaking, 'adaptive behavior' means behavior which is adjusted to environmental conditions."

22 - "Despite the ubiquity of such responses as reflexes, taxes and fixed-action patterns, animal behavior is by no means solely reactive. Factors internal to an animal can also play an important role in the initiation, maintenance, or modulation of a given behavior.... Behaviors which show no simple or rigid dependence on external stimuli, but are instead governed primarily by the internal state of the animal, are known as MOTIVATED BEHAVIORS. In these behaviors, an animal's propensity to exhibit a given behavior such as feeding depends not only upon the presence of the appropriate environmental stimuli (i.e. food), but also upon typically MOTIVATIONAL variables (i.e. hunger)."

23 - "Any individual animal consists of a large collection of reflexes, taxes, and fixed-action patterns, many aspects of which are under at least some motivational control. As an animal confronts its environment with this diverse behavioral repertoire, it must properly coordinate its many possible actions into coherent behavior directed toward its long-term survival. Toward this end, the behavioral repertoire of a natural animal typically exhibits a certain organization. Some behaviors normally take precedence over others. Some behaviors are mutually exclusionary (i.e. any behaviors which utilize the same motor apparatus for incompatible actions). Switches between different behaviors depend both upon environmental conditions and internal state. These relationships are often described as rigid and strictly hierarchical, with cleanly delineated behaviors and simple all or nothing switching between them. In reality, the relationships may be nonhierarchical, the organization can change depending upon the behavioral context, and the behaviors can partially overlap so that discrete switches between them are sometimes difficult to identify."

25 - "In particular, I strongly believe that the behavior of simpler animals has all of the ingredients which artificial autonomous agents require in order to flexibly cope with the real world: it is goal-oriented, adaptive, opportunistic, plastic, and robust. While the specifics of any given animal behavior are unlikely to be of direct use to an engineered agent, the general principles most certainly are."

25 - "Consider the following problem: You must design the control system for a device which can autonomously accomplish some open-ended task (such as 'stay out of trouble' or 'keep this area clean') in a complex, dynamic, unpredictable, and, in many ways, openly hostile environment. You have considerable general information about the structure of this environment, but

cannot assume that this information is complete in any sense. Your system must therefore be capable of flexibly applying whatever behavioral repertoire you choose to give it to the actual situations it encounters."

71 - "Locomotion is... an interesting adaptive behavior in its own right. An insect robustly solves this complex coordination problem in real time in the presence of variations in load and terrain, developmental changes, and damage to its walking apparatus.... Because this [model] insect can fall down, its locomotion controller must properly coordinate the movements of the insect's six legs in order to produce successful walking. The insect must also be able to walk at a variety of different speeds while maintaining the stability of its body. In addition, the locomotion controller must be robust enough that small perturbations (such as those caused by a collision with a brick or a wall) will not seriously disrupt it."

125 - "Feeding is the first example of a goal-oriented behavior in P. COMPUTATRIX."

127 - "Feeding is a prototypical motivated behavior in which attainment of the goal object (food) is clearly crucial to an animal's survival. In this case, the relevant motivational state is hunger. When an animal is hungry, it will exhibit a sequence of APPETITIVE BEHAVIORS which serve to identify and properly orient the animal to food. Once food is available, CONSUMMATORY BEHAVIORS are generated to ingest it."

139-140 - "... there is no simple neural correlate to the artificial insect's 'desire' for food. Indeed, the so-called feeding arousal [FA] neuron appears to be something of a misnomer. From a neural perspective, the insect's response to food is the result of a complex dynamics of interaction between an internal positive feedback loop (mediated by FA) and a negative feedback loop (mediated by [neuron] ES) which is closed through the external environment."

164 - "P. COMPUTATRIX is clearly a reactive agent in the sense that its behavior is quite responsive to any contingencies which arise in its interactions with its environment. An insect which is tracking an odor signal to food, for example, switches immediately to edge-following when it encounters an obstacle in its path. However, the artificial insect is not purely reactive, because it is also capable of organizing its individual behaviors in such a way as to achieve particular objectives. An insect which is hungry generates very different sequences of behavior than one which is not... this goal-oriented behavior... derives from the ability of certain internal states of the insect's nervous system (e.g. those resulting from the interacting feedback loops underlying its feeding arousal) to modulate the interactions between the neural circuits responsible for its various behaviors."

>I don't see the value of having detailed, realistic models of a phenomenon
>that does not exist.

We think Beer is modeling a phenomenon which DOES exist: behavior, the control of perception.

>I think the reason for the greater interest in Beer than Bill is the
>the reason for the greater interest in fuzzy logic, expert systems, non-linear
>systems, etc than PCT -- it's because most students of living systems
>assume that organisms generate output rather than control input. They
>assume that behavior is the kind of phenomenon that is consistent with their
>models and methodologies. This is why there are probably no more than 5
>psychologists in the country who work in the context of PCT.

We disagree. Beer's model achieves invariable ends by variable means. THAT is
impressive!

>If knowledge were a popularity contest, I'd jump on the Beer bandwagon
>immediately.

We wouldn't advise it.

>What Tom meant is that PCT tries explicitly to model CONTROL. Beer's models
>do not.

We think that they implicitly (and maybe even explicitly, in some cases) model
CONTROL.

>This is THE big difference to me. The fact that some control emerges
>out of the Beer model is nice; but if that was not the goal of modelling then
>I don't see how you can evaluate the model (except by continuing to
>get surprise results).

How could it be a surprise? Nobody has been able to model "variable means to
invariable ends" any other way, legitimately. And Beer didn't act surprised.
He built in control, because he had to -- that is what you should be telling
the world!

>Also, if the "neuro-physiologically correct" components do not produce a
>system that controls the variables that it is known to control then who is
>right?

What?

>If the bug controls the angular position of its leg (for example) and we find
>no circuits that allow this kind of control then do we say that the bug isn't
>really controlling that variable?

What?

>It seems to be that, by going with the "neuro-physiology" we could end up
>like the (possibly apocryphal) engineers who found that it was impossible for
>bumble-bees to fly.

It seems to us that, by going with ONLY Bill's detailed model, we could end up
claiming that a controlling model doesn't control!

>I would like to know more about the control abilities of the simple organism's
>whose CNSs are so well known (if they really are).

We can provide plenty of references.

>Maybe this data is available but I saw nothing in Beer's book that suggested
>that he was trying to design into P. computatrix the ability to control any
>variables.

See above quotes.

>It's not the model I cling to -- it's the phenomenon. The fact that we can
>model the phenomenon gives confidence that it can be produced by a real
>system -- this takes the "metaphysical mystique" from the phenomenon of
>purpose (control) which many scientists once rejected as fancy. Bill and

>I have built models where higher order systems alter properties of lower order
>output functions rather than their references (this was not part of
>the original model) because it looked like this is what happens. I am
>happy to use Beer-type components to model control. But that's what I
>would want to use them for -- to produce control, not just some "output"
>(like a "gait") that happens to catch one's fancy. First, I want to
>know what variables are involved in "gait" and whether they are controlled.
>Then I would start to model.

Still building bridges, eh? We fully support such endeavors as are described above.

>What would constitute a test of Beer's model?

Agreement with behavioral and physiological data.

>Is Beer's model successful if the components behave correctly even if the
>simulated bug can't balance on uneven terrain?

Disregarding the fact that Beer's bug is a 2-D model, it would not be successful with regard to balancing on uneven terrain, but the components could behave successfully in other ways -- so the modeler has to make revisions to meet the failures as they appear.

>Beer is not TRYING to model control.

See above.

>The "contrary possibilities" seem to be in the building blocks. If one of
>the "contrary possibilities" is that organisms generate outputs, then we know
>already that they do -- these are accidental side-effects of control.

That isn't what we meant. We meant CONTROL structures contrary to Bill's detailed model.

>But don't forget those BEHAVIORAL CONSTRAINTS. That is what I think is
>seriously missing from the Beer work.

We agree! Time to get going on that...

>He may think he is working under behavioral constraints, but his concept of
>behavior is very different than mine.

From the foregoing, we hope you'll understand why we think that the "very" seems exaggerated.

>Again, I am not interested in bashing Beer. He has done lots of very
>interesting work. I just don't immediately see its relevance to my
>interest in purposive behavior. But maybe you (or he) could help clear that
>up for me.

We've been trying! Once more: Beer has a CONTROLLING model which is generating a good deal of interest. PCT could introduce itself to a lot of folks via a paper pointing out that the theory underlying the wonderfulness of Beer's bug (namely, its robustness) is PCT.

Again, CUTE pun!

Pat & Greg

Date: Tue Feb 18, 1992 5:25 pm PST
Subject: LFG system re-done

I've now redone the system as `self-extracting' archives. So, the new instructions would be:

anon ftp to `durras@anu.edu.au'

collect aaa_read.me plus the files lfg*.exe

transfer at least lfg-pc1.exe and lfg-pc2.exe to your PC, using binary file transfer settings, and unpack as described in aaa_read.me. (You might or might not want the other two archives - they're presently there for the benefit of some other people who wanted to look at them.

NOTA BENE: this is *not* a PCT model of anything, but does, I believe, actually encode quite a lot of what linguists presently know about language. Working through the exercises, as described in exercises.doc, might be the best way to approach it, tho they weren't designed to function as a free-standing teach-yourself course.

Avery.Andrews@anu.edu.au

Date: Tue Feb 18, 1992 5:29 pm PST
Subject: language

re Powers (920218.0900)

>>>How does the program know that "bite" takes arguments of classes SUBJ and
>>>OBJ, whereas "dog" and "cat" don't?
>>
>>This is specified in the lexicon ...
>
>OK, this takes us another step. What are the processes that build the
>lexicon? If you do it yourself, how do you know that "bite" is of a class
>that takes a SUBJ and OBJ class of arguments?

You do it yourself, and learning how is part of becoming a linguist.

Figuring out how language learners actually do it is another question, without a quick & satisfying answer. One of Chomsky's main points has been to insist that this should be seen as the central question of linguistics, though many people aren't very happy with his proposed answer, which is that there is rather highly constrained class of attainable grammars (determined by our genes), plus some device for sifting through them to find the one that `best fits' the linguistic behavior of the community that the learner (generally assumed to be a child) finds themself in.

A story about how this might work for LFG is told in

Pinker (1984) "Language Learnability and Language Development", MIT Press

Though Chomsky & his closest followers would go for something rather different these days (`parameter-setting' in the Government-Binding theory).

Avery.Andrews@anu.edu.au

Date: Tue Feb 18, 1992 5:39 pm PST
Subject: beer bug whinge

I too have a bit of a gripe about the bug debate, along the same general lines as P&GW, which is this.

When as a linguistics student, I used to get into arguments with psychology students (e.g., Skinnerians), the thing that used to drive me batty about these people was that they seemed completely oblivious to the fact that they didn't have an `effective description' of languages, & couldn't possibly get a half-decent one on the basis of the ideas they were using, while we were at least some way down the track towards having such a thing, an `effective description' being one that can in some sense be implemented so as to enumerate or detect sentences of the language.

An `effective description' is not as nice a thing to have as `generative model' in Maturana's sense (as I understand it from reading posts here), but it's still a lot better than the mess of verbiage & statistics that the psychologists seemed to be serving up. What makes it nice is that it cranks out quite detailed predications about what strings should be OK sentences and what they should mean, as a matter of objective and essentially mathematical fact (though the math is shallow and boring). So wrong ideas in due course manifest their wrongness, and people who want to work on the mechanisms underlying the behavior can get detailed guidance on what these mechanisms should be able to do.

So it seems to me that beer's bug ought to be greeted with open arms by anyone interested in understanding behavior, because the thing does indeed behave, on the basis of what people have been saying about it (I confess to be waiting to get some of my own projects finished, and for the next NSCK version, before getting stuck into it myself). So what if a lot of it is made up, or doesn't look like what you would expect on PCT principles, or fudges the physics -- all this sort of thing can be sorted out in the passage of time. I think there's a vast chasm between models that are sufficiently specified to produce some sort of behavior, and those that aren't, and that at the outset it REALLY DOESNT MATTER if the behavior is wrong in various details and produced in the wrong way.

People on the Beer/Chomsky/PCT side of this chasm can convince each other of things by producing models that are *obviously better* by virtue of covering the data more accurately, being in closer conformity to what is know about how the nervous system works, etc. In principle, that is, though in practice its not so easy, because alternative models (at least in linguistics) tend to do better in different areas, so that people have to make choices on the basis of what they judge to be the promise rather than the actual current performance. On the other side of the chasm, who knows (and who cares)?

  Avery.Andrews@anu.edu.au

Date: Tue Feb 18, 1992 6:17 pm PST
Subject: BEERBUG debate finis

From Pat & Greg Williams (920218-3)

>Bill Powers (920218.0900)

>I have a better idea. Why don't we abandon this stupid argument? You do
>what you're interested in, I'll do what I'm interested in, we'll
>collaborate where the domains intersect, what will come of it will come of
>it, and we'll still be friends.

Done. If you want us to answer anything else on this topic, let us know. We are not trying to avoid your questions, but additional debate appears fruitless, if not stupid.

Best wishes,   Pat & Greg

Date: Wed Feb 19, 1992 7:46 am PST
Subject: W.T.P. - Paper

The "Materials related to living control systems theory" by Greg Williams contain the citation of this article (by W.T. Powers 1986 (?)):

     Control theory, constructivism, and autopoiesis

     unpublished paper (11 pages).

Has anybody out there ever seen this paper? I am very interested to get it. Can anybody help me?

Wolfgang Zocher

email: Wolfgang.Zocher@cdc2.rrzn.uni-hannover.dbp.de

Date: Wed Feb 19, 1992 9:48 am PST
Subject: LFG system re-done (fwd)

>From: Avery Andrews <andaling@FAC.ANU.EDU.AU>
>Subject: LFG system re-done
>
>I've now redone the system as `self-extracting' archives. So,
>the new instructions would be:
>
> anon ftp to `durras@anu.edu.au'

I can't reach this host. If anyone wants to upload athe files to biome, I'll move them into the csg area so that they will be available on this side of the world.

Bill

--
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: Wed Feb 19, 1992 10:02 am PST
Subject: LFG system on BIOME

Ignore my previous message. I carelessly copied a typo:
> anon ftp to `durras@anu.edu.au'
should read:
> anon ftp to `durras.anu.edu.au'
>
> collect aaa_read.me plus the files lfg*.exe

Anyway I'm logged on right now and am transferring the files to biome, so that they will be available by anonymous ftp from here too, and, more important, they can be fetched from the mail server.

When I get the files together I'll probably put them in a separate directory. Wait until tomorrow and then get the Index from csg for details (send message "get csg/Index" to biome.bio.ns.ca).

Bill
--
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: Wed Feb 19, 1992 1:17 pm PST
Subject: beer bug and onward

[From Rick Marken(920219)]

I'm as ready to drop the Beer Bug topic as anyone (at least until Beer himself joins in -- if he does). I will try to segue into a new (but related) topic in my response to the following comment from Avery Andrews:

>So it seems to me that beer's bug ought to be greeted with open arms
>by anyone interested in understanding behavior, because the thing
>does indeed behave

PCT models behave, too. In fact, they control. And they mimic human control (as Bill mentioned) to within 5 per cent (far better in most cases). These are not statistical matches -- the model produces behavior that matches human behavior perfectly every time. So having a model that "behaves" is not, per se, something that makes PCT people sit up and take notice. As I mentioned earlier, automata have been designed since the 16th century that produce interesting behavior -- have you ever watched some of those german clock automata do their thing? There are modern industrial robots that really go through their "behavioral" paces. The problem with most of these automata is that they cannot behave in real environments -- where the results of their actions are influenced by random disturbance as well as by their own efforts. Thus, industrial robots must work in fairly well "controlled" environments -- where potential disturbances are held constant. There is some ability to deal with disturbances (they can adjust to slight discepencies) -- but, by and large, they are fancy "action producers". The interest of Beer's bug is that is does seem to control -- that is, produce consistent behavioral results in the context of certain kinds of disturbance. Those of us who have greeted Beer's bug with something less than open arms have been reluctant mainly because Beer did not seem to recognize the basic organizing principle (control of perception) that gives the bug what "adaptive" capabilities it has. Pat and Greg have argued that this failing is far outweighed by the virtues of the physiological realism of the bugs infrastructure. I think it is at this point that we've just agreed to disagree.

So here is where I propose a new, but related, topic. I think that a great deal of the attraction of Beer's bug is the visual appearance of a bug on the screen searching for food, finding it, eating and then looking again. I would like suggestions about a similar project that would illustrate the strengths of building a behaving system from a PCT perspective. I think the "little man" and "gatherings" demos do this already. I like to think the "random walk" bug does it too. But maybe, to bring the point home, we need something visually compelling (like the bug) such as a "juggling" person, a ballet dancer or maybe a skier -- I think it should be an agent whose behavior requires control of relatively complex percpetual variables (this is part of the problem with the PCT demos so far -- the variables controlled are relatively "simple"). Also, I think that the main source of disturbances should be the actions of the agent itself; it's just too tough to model a realistic external environment that will affect all your controlled variables. Also, the effect of the agent on the controlled variable should be relatively inconsistent -- another source of "disturbance". Thus, if you build a juggler then the outputs that move the arm up to throw the ball should always have slightly different effects on the ball.

I would be willing to work on a "behaving agent" program. But I would like to have suggestions about what might be the "best" kind of program to work on -- in the sense that it might be seen by people (PCTers and non PCTers) as a demonstration of the value of the PCT approach. Even if I don't actually make the program, the discussion might reveal some of the kinds of things that people want to perceive about a "successful theory of behavior"

Best regards

Rick

Date: Wed Feb 19, 1992 1:29 pm PST
Subject: BEERBUG; Andrews' LFG

[From Bill Powers (920219.0900)]

G&P Williams (920218-3) --

Re squabbles: Good.

I'm still waiting for an ILL copy of the Beer book, so we've finally broken down and ordered it.

A suggestion for the next version of NSCK, if it's not too late:

Provide for antenna touch sensor response proportional to distance from obstacle over some short range (like 10% of antenna length). This is equivalent to letting the antenna bend and sensing the amount of bend. Provide a sensor for each direction of bend and for longitudinal (spearing) contact. A standard Beer neuron can convert the proportional effect back to on-off by adjustment of gain and threshold, so this simply adds a degree of freedom to the model. Distinguishing directions of antenna bend will help in getting out of corners.

Can a cockroach move its antennae relative to its head?

Avery Andrews (920218.1729) --

>>What are the processes that build the lexicon? If you do it yourself, how
>>do you know that "bite" is of a class that takes a SUBJ and OBJ class of
>>arguments?

>You do it yourself, and learning how is part of becoming a linguist.

But if this is to be a model of how human beings use language, it has to work even if they're not linguists. The language structures depend on existence of a lexicon. The lexicon must exist in every person who successfully uses those language structures, because the lexicon contains a great deal of information without which the rest of the processes can't work. That's why I ask you to examine the processes by which you construct a lexicon -- by which anyone would, linguist or not. If we can explicate those processes in terms of more basic operations, we'll have another clue about how to represent language in the HCT model. So far I see class-recognition perceptual processes, sequence perception, and perception of logical conditions. Below that we have to have some way of deciding what class a word like "bite" belongs to, and I think that way probably has something to do with lower levels of perception like relationship, event, transition, configuration, sensation, and intensity -- as nonverbal perceptions, not words.

If ONLY linguists do these things, doesn't that raise a question as to what this game has to do with natural language?

Best to all

Bill P.

Date: Wed Feb 19, 1992 2:01 pm PST
Subject: connecting to durras

Try connecting to the IP address 150.203.22.8 rather than the host name durras.anu.edu.au. When I gave Avery a file recently, I had to do even this from a local Sun server, my mail host (an elderly C/70) was not able to make the connection. Don't know why it timed out, but it did. So there may be other problems, but this at least gets around the problem of your host not having up-to-date host tables.

I'm going to have to be a quiet bystander for another week or so. But I do want to observe that Bill and Greg/Pat have just given us a very mild version of the sort of quote-riposte exchange that swallows up all else in the usenet world. I'm glad there seems to be consensus not to open ourselves up to that.

Bruce

Date: Wed Feb 19, 1992 6:49 pm PST
Subject: language acquisition

Re Bill on grammar & lexicon acquisition:

In Chomsky's, view, which I'd pretty much go along with, one can:

(a) try to get a better idea of what mature languge is like

(b) study acquisition, based on some provisional answer to (a)

What doesn't make sense is to study acquisition `on its own', as it were, without any attention to the nature of what is being acquired.

So people working in the area simply have to make a bet on when they have done enough w.r.t. (a) to do anything sensible w.r.t. (b), and I think it really is a gamble. The Pinker book I mentioned in an earlier posting develops an account of how LFGs might be acquired, but the framework itself also needs a lot more work -- there are all sorts of grammatical structures that it doesn't give a very good account of.

My current general thinking on grammar acquisition is:

1. grammatical structures have a fairly tightly defined universal format, as well as a highly constrained relationship to meaning. These are both genetically specified, but *not necessarily as some sort of `autonomous' language faculty*. There is also a strong innate component to meanings, especially at the level talk about food, furniture, animals & other things that are central to the experience of young children.

2. one of people goals is to perceive those around them to be saying things that make sense.

3 . So, upon hearing a string of noise, one tries to find a grammatical structure, and meanings for the words in it, that will make it a sensible thing to say. E.g. if mommy is standing in front of you, holding a cookie, and says, in a loving tone of voice:

    "do you want a cookie"

   It is not sensible to guess that this means `that gorilla is about to rip your arm off'.

4. Grammatical rules are then abstracted from structures on the basis that (a) they permit the structures that are heard (b) don't generate too much that isn't heard. ((b) is actually rather problematic).

The goal of perceiving other people to be making sense would also play a substantial role in ordinary language comprehension, where it motivates our ignoring contextually implausible meanings. E.g. if I say `I'm need to go the bank to get some money', you don't think of the river.

But any substantive fleshing out of this picture, whether at the rather abstract/cognitive level pursued by Pinker, or something with a more neurological orientation, would depend a lot on what you thought grammars were in general like. I.e. the shape of the space of possibilities that the language learner is navigating through in his search for the local language.

Avery.Andrews@anu.edu.au

Date: Wed Feb 19, 1992 7:27 pm PST
Subject: Segue

From Pat & Greg Williams (920219)

We hope this doesn't count as squabble reborn, but as asking for a point of fact, which appears to be allowed in the transition to non-squabbling.

Question: Where does the 5% (or better) accuracy figure for PCT models come from? Which models, what data? I (Greg) admit to surprise regarding this figure for the Little Man, as I didn't know Bill had compared it to "real" data.

Thanks,

Pat & Greg

Date: Thu Feb 20, 1992 11:50 am PST
From: Dag Forssell / MCI ID: 474-2580

TO: csg (Ems)
EMS: INTERNET / MCI ID: 376-5414
MBX: CSG-L@VMD.CSO.UIUC.EDU
Subject: Beer debate
Message-Id: 71920220195017/0004742580NA3EM

From Dag Forssell (920220)

Reference: Various

I must write to express my appreciation for the Beer debate. From my perspective, it has been very useful. When the dust settles (all the error signals have caused appropriate reorganization), the CSG group will have made progress.

The strength of CSGnet is the open, serious debate on global as well as detailed issues. As a listener, I have enjoyed the debates on epistemology and social control, and have gained clarity on PCT through them and others.

Another strength is in the personal courtesy, which is a prerequisite for anyone to spell out in detail an understanding and ask for comments.

One difficulty and at the same time "stimulus" for the group is the diverse membership. Backgrounds and PARADIGMS OF THINKING vary from the liberal arts to theoretical physics.

Personally, I graduated in 1965 with an MS in mechanical engineering, but have not emphasized equations much since. However, I can recognize that it affects your paradigms of thinking in everything you do, read or ponder throughout your life.

Rick mentioned "Flat Earth'ers" (in connection with behaviorists - what else), and how they, too have an explanation for why ships (seem to) disappear over the horizon. That was interesting to me, so I asked Rick directly. He questioned his source, and the answer came back that due to the greater density of the air at sea level, compared with the lesser density a few feet up, the light bends upward as it travels away from the ship. No flat earth person will bother to ask: How much?

I recall reading about Perpetuum Mobile - perhaps in my teens. Even today, people come up with clever mechanisms that will give more energy (perhaps from weights falling, away from the center of a wheel) than they consume (perhaps from the same weights being raised again on the other side of the same wheel, but closer to the center, with less leverage on the wheel's rotation).

The books I vaguely remember show dozens of these PM machines built in the 1700's and 1800's, complete with intriguing pictures.

Bill Powers told me in Durango last August how he as a teen took two electric fans and placed them facing each other. He spun one with his finger to generate electricity, allowing its motor to act as a dynamo. Connecting the two power cords to each other, Bill expected the wind created by the second fan to power the first, generating still more current. Bill's experiment did not work very well.

I am telling these stories to illustrate the timeless appeal of qualitative thinking, unchecked by quantitative considerations.

A liberal arts education, including psychology (I don't know this, but it is my paradigm at the moment), does not teach or encourage quantitative thinking.

With qualitative thinking, all you can possibly do is to observe phenomena and create some statistics.

With quantitative thinking, you can correct qualitative mistakes and go into much farther depth.

I shall insert here my posts (Forssell 920121), which were an attempt to clarify this point of view:

Subject: How is PCT different? Take two.

(Recognition due to Greg Williams for n-1 explanation idea).

PCT is a paradigm or theory of human activity which offers you greater powers of explanation than any other approach I know of.

Let us talk about theories.

For many people, a theory is nothing more than a generalization of some observed fact. As an example, I might suggest: I have a theory that all Swedes are blond. (All the ones I have ever seen in the movies sure look that way). You might disagree and say that only 80% of them are blond. Some will argue that if you can find one who is not blond, the theory is refuted.

Some scientists argue over percentages of observed facts. These are sometimes called "soft" scientists.

Other scientists do not think that percentages of observed facts without explanation qualifies as science, but want explanations, what they think of as theories, for the observations.

Think about the "hard" sciences. If you ask about a mechanical phenomenon (a fact you have observed), a scientist can provide an explanation. If you ask him about the explanation, he can explain it. In the hard sciences, these explanations or theories go many layers deep. Explanation for explanation for explanation for explanation and so on. As they have been suggested over time, these explanations have been used to predict what will happen under many varied circumstances and when they appear to predict without ever missing, the scientist thinks of his explanations as laws of nature. (The explanations seem infallible).

When you dig very deep, you get down to the atoms and parts of atoms. You learn explanations for forces within atoms as well. When you press the scientist, he will admit that our explanations at this level are guesses only. There is no agreement on these phenomena. Interestingly, you find that chemists, biologists, physicists, electrical engineers - you name them - are all looking at the same explanations at the bottom of the explanation "layer cake."

What is different about PCT is that instead of being satisfied to argue statistics about phenomena of the mind (without any explanation at all), PCT provides several layers of explanations in a detailed model. The model is used to provide predictions (under many varied circumstances) which can be tested against experience and give you confidence.

PCT in its approach to theories and explanations is more a hard science than a soft one. That explains why it is much more useful, but also why it requires more effort to study it in depth.

This can perhaps be illustrated as an upside down layer cake "pyramid" of explanations/theories.

Mechanical Phenomenon	Electrical Phenomenon	Chemical Phenomenon	Biological Phenomenon	Living Phenomenon

```
\  Explanat   \  Explanat  |  Explanat  |  Explanat /  Explanat     /1
  \____________\___________|____________|__________/______________/
    \    explanation for expl  |  exp f exp  |     exp f exp    /    2
      \________________________|_____________|________________/
        \    explanation f exp f exp  |  exp f exp f exp    /       3
          \___________________________|___________________/
            \       explanation f exp f exp f exp       /           4
              \_______________________________________/
                \         Fifth level deep          /               5
                  \_______________________________/
        |           \         Atoms           /       ^             6
        |             \_______________________/       |
        |               \     Protons       /         |             7
        |                 \_______________/           |
        v                   \  Quarks   /             |             8
   EXPLANATION                \_______/           PREDICTION
                                \ ? /                               9
                                  \/
```

This is intended as a conceptual chart. I have no basis for how many levels to suggest.

Recent discussions on the net have been very helpful to me in sorting out the Popper vs Kuhn vs Powers vs Statistics arguments in my mind.

My personal committment is to teach applications of PCT to industry. I will not take the time to read Beer at this time. I am hard at work on the third version of my program. (You could perhaps call it: How to teach quantitative thinking in a qualitative world). I have learned to address the phenomenon first, then the application, third the theory - only the global model. The Beer debate has been helpful to me in that I will be very cautious with the detailed model. I will not need controversy with that part. It would detract from teaching the global model.

I do want to take the time to attempt to contribute to the Beer debate in this way.

Without a solid appreciation of quantitative thinking, it is hard to appreciate Greg's point of view. This makes it hard for Greg to be understood, regardless of the merits of his content.

I will now take issue with Rick's post (Rick 920219):

>Those of us who have greeted Beer's bug with something less
>than open arms have been reluctant mainly because Beer did not SEEM TO
>recognize the basic organizing principle (control of perception) that
>gives the bug what "adaptive" capabilities it has. Pat and Greg have
>argued that THIS FAILING is far outweighed by the virtues of the
>physiological realism of the bugs infrastructure. I think it is at this
>point that WE'VE JUST AGREED TO DISAGREE.

This is not what happened as I see it. There is an alternate choice of perception available. It is not fair to offer a questionable summary of Greg and Pat's effort after they have agreed to shut up. This creates a strong error signal in me.

What happened as I see it, is that the two giants of quantitative thinking in CSG had a sincere exchange about their perceptions on the merits of the details of PCT.

Greg and Pat have exhibited great initiative, courage, clarity, writing skill and deep knowledge in initiating this thread. (Driven by an error signal from being unsuccessful in interesting CSG in Beer's work).

Bill (as usual) demonstrated his great openness to different information and answered in as much (plenty) detail as he is capable of on short notice.

Error signals arose at both ends as must be expected, where personal understandings must be adjusted if a common one arrived at.

It is my perception that Greg and Bill are very close to an understanding.

Greg and Pat have consistently argued, most clearly in their post (Greg & Pat 920218), not that

> .... because Beer did not SEEM TO
>RECOGNIZE THE BASIC ORGANIZING PRINCIPLE (CONTROL OF PERCEPTION) that
>gives the bug what "adaptive" capabilities it has. PAT AND GREG HAVE
>ARGUED THAT THIS FAILING is far outweighed............

but that Beer deliberately and explicitly models CONTROL. Beer's sin is not to use PCT terminology.

Let us at least read what Greg and Pat are writing.

As I read Bill's "better suggestion" and Greg's reply, the end of the debate (for now) was strictly a matter of personal priorities. There was no hint of "agree to disagree" as I read it.

As tempers cool, insight grows on all sides and the relevance of Beer's work emerges, the debate will and should re-emerge as well.

Best to all,

Dag Forssell
23903 Via Flamenco
Valencia, Ca 91355-2808
Phone (805) 254-1195   Fax (805) 254-7956
Internet: 0004742580@MCIMAIL.COM

Date: Thu Feb 20, 1992 1:39 pm PST
Subject: PCT as dogma

RE: The (now historical) PCT as dogma discussions:

#ifdef You_Care_About_My_$0.02

When I first heard of Beer's work via Pat & Greg's NCSK,
I ran to the local UC to get a copy of his book and read
it. When I later got a pointer to CSGnet, I again ran to
the local UC and got a copy of Power's book -- and am

currently working my way through it. This is how I gather facts and make informed decisions. Yet on this net I have repeatedly heard LOUD objections to Beer's work which usually ended with words to the effect "... I haven't read his book or studied the NCSK code, but there is obviously something wrong with it..."

Also, besides Beer and having originated about the same time is some work by Hugo De Garis of CADEPS Artificial Intelligence and Artificial Life Research Unit, Universite Libre de Bruxelles, Brussels, Belgium. He started with a controller for basic leg motion, went on to evolve a controller for bipedal motion (called the walker) and then went on to develop LIZZY - a quadroped lizard who searches out mates, food and runs from predators. He is using what he calls GenNets to evolve the various controllers in his hierarchy.

His words: "The success of walker had such a powerful effect on the author that he became convinced that it would be possible to use the basic idea that one can evolve behaviors (the essential idea of genetic programming) to build artificial nervous systems by evolving many GenNet behaviors (one behavior per GenNet) and then switching them on and off using control GenNets. ... One final idea concerns the possibility that elaboration of the middle "logic" layer (i.e. between the detectors and the effectors) in artificial nervous systems may become sophisticated enough to be capable of handling "symbolic reasoning"..."
From: Neural and Intelligent Systems Integration, Branko Soucek and the IRIS Group, 1991 John Wiley & Sons Inc.

See Also:

[1] H. de Garis, "Genetic Programming: Modular Neural Evolution for Darwin Machines, "Proceedings International Joint Conference on Neural Networks, Jan. 1990, Vol 1, pp 194-197.

[2] H. de Garis, "Genetic Programming: Building Nanobrains with Genetically Programmed Neural Network Modules," Proceedings International Joint Conference on Neural Networks, Vol. 3, June 1990, pp. 511-516.

[3] H. de Garis, "Genetic Programming: Building Artificial Nervous Systems Using Genetically Programmed Neural Network Modules," Proceedings 7th International Conference on Machine Learning, 1990, pp 132-139.

What is fascinating about this work, is the notion of building complex behaviors from the interaction of simpler ones.(The global/local duality again).

For those who find fault with Beer`s use of present neurophysiological data have heart: de Garis' neural modules have NO neurophysiological basis whatsoever. Brute force attacks using totally connected recurrent networks with no basis in reality and no underlying theory, you see, work too. The question I ask myself is: "Which approach will take me the furthest?"

I for one find reality hard to dismiss. But, am currently in need of an underlying theory from which to base my models of intelligence as emergent behavior. Using neurophysiological data alone will only take one so far. What do you do when the data runs out? Beer's approach was to guess at a possible neural configuration. De Garis' approach was to guess from the very beginning. Does PCT provide a theoretical basis from which one can do better than guessing when the data runs out? Will it allow a basis from which to launch a search of the space of possible control configurations for the development of intelligent agents which exhibit particular behaviors?. Not if open-minded, scientific investigation is replaced with dogmatic attacks of the form featured here.

Sermon over. Back to science as usual.

#endif

Alan E. Scrivner c3141aes@mercury.nwac.sea06.navy.mil

Date: Thu Feb 20, 1992 2:14 pm PST
Subject: Segue

[From Rick Marken (920220)]

Pat & Greg Williams (920219) ask:

>Question: Where does the 5% (or better) accuracy figure for PCT models come
>from? Which models, what data?

Tom Bourbon is the one to answer this. Maybe he's locked out of the net right now so I'll give it a shot. My 5% accuracy comment was motivated by my experience with very simple models doing tracking tasks. One example is the two level model described in my JEP:HPP article. The model controlled the distance between pairs of lines and the relationship between these distances. The position of the lines was influenced by the value of a handle controller (for the subject) or a value generated by the model (for the model). In both cases line position was also influenced by a slowly varying random disturbance. I measured accuracy of the model as the correlation between values produced by the handle (subject) with those produced by the model in the same environmental (disturbance) conditions. The correlations (based on 400 pairs of data values) were always in excess of 0.98. This means that the model accounts for 96% of the variance in subject response. Tom Bourbon has done even better -- as I recall -- accounting for over 98% of the variance in his data. It's not just the accuracy that is impressive (to me) but the consistency. The model matches behavior at this level every time (for well practiced subejcts -- those who can control well) and in novel circumstances (like with new disturbances and, in one of my favorites, when you change the polarity of the connection between handle and cursor-- during the "positive feedback" phase the model continues to matche the subject perfectly (r = .95 or so) with no change in the model). Bill Powers has developed the control model so that it not only matched subject output behavior but also matched the behavior of the controlled variable -- which is amazing because when the controlled variable is held in a fixed reference state, variations in its value seem almost random -- but they are not. I forget what kind of accuracies Bill got when predicting controlled variable behavior. But it was quite impressive to me -- though perhaps not at the 5% or better level.

Regards

Rick

Date: Thu Feb 20, 1992 2:16 pm PST
Subject: Language; models; 5%; new book

[From Bill Powers (920220.0900)]

Rick Marken (920219) --

>So here is where I propose a new, but related, topic. I think that a great
>deal of the attraction of Beer's bug is the visual appearance of a bug
>on the screen searching for food, finding it, eating and then looking
>again. I would like suggestions about a similar project that would
>illustrate the strengths of building a behaving system from a PCT
>perspective.

This is a good idea. I think, however, that to make it more feasible we should give up on complex things like skiing and juggling. The reason is not that the control systems would be too hard to design, but that simulating the physics would be too difficult. Beer uses a very simple two-dimensional environment, and that seems like a good idea to me, too.

Also, I suggest that we not try to model the control systems in Beer neurons right away, but build a block-diagram-level system first. Once the block-diagram model works, we can either propose neural circuitry (using NSCK) that would accomplish what is in the boxes, or leave that to the neural modelers to figure out.

In defining the behavior of models in simulated environments, it's important not to use the modeler's perceptions in specifying what the system is to "do." The model would not "search for food" -- that is an observer's impression. The model would try to bring some sensory indication of the presence of food to a reference level. Its actions might look to a human observer like "searching" and the object might appear to be "food," but unless the model itself contains such notions, this can't be a correct description from the model's point of view. Even "edge following" is a human interpretation -- the cockroach doesn't perceive anything corresponding to "edgeness" or "following." It just controls its antenna touch signal. By doing so it appears to us to "follow an edge," "navigate around an obstacle," and so on. That is how a human onlooker interprets the results, anthropomorphizing.

What we're trying to do, I think, is to construct a system that controls for various perceptual variables by use of its actuators. In a given environment, this will lead to much behavior that the human observer can see. The human being will see all sorts of high-level variables in the movements of the model, and marvel at the "intelligence" demonstrated by it. This happens all the time with the Gatherings program. People will say "Hey, that was a dumb move, he could have turned right and taken a much shorter path!" It's hard to explain that the actors in this program don't know anything about paths or shortness or obstacles or goal-circles. They're just controlling sensed proximity, left and right, without knowing WHAT is proximate to them, and in fact without knowing anything about space or speed or direction. Someone who knows the controlled variables will object to such observations: "No, he's just trying to keep proximities under control by varying whatever affects them." But that's hard to

understand unless you're used to taking the point of view of the model, with its limited perceptions, instead of the observer -- who can see, actually, too much.
If you define an organism that can move and turn, we have in the Beer model one example of such an organism, so can just take walking and turning as output functions that can be used in higher-level control systems. There's no need to reproduce the whole circuit here; all we need to know is that one output signal can vary forward speed, and another can vary direction of movement. The speed signal can come from a comparison of the sensed strength of a sensory input with a reference strength, and the direction signal from the difference between two sensory signals compared with a reference-difference. If you want a more complex organism, you can provide sensitivity to different substances: food, water, pheremones, light, air velocity, temperature, and so on. The environment could be designed to provide these things in varying amounts in different places. The reference levels for these things could be arranged to depend on variable inner states, as Beer has done with hunger. Of course the more of these variables you put under control, and the more complex their arrangement in the environment, the more complex the actions of the system are going to look, despite the underlying simplicity. The point of the control-system model is to show that you can get a lot of apparent complexity of behavior out of controlling just a few simple perceptions in a complex environment.

Avery Andrews (920219) --

>My current general thinking on grammar acquisition is:
>
> 1. grammatical structures have a fairly tightly defined universal
>   format, as well as a highly constrained relationship to meaning.
>   These are both genetically specified, but *not necessarily as
>   some sort of `autonomous' language faculty*. There is also
>   a strong innate component to meanings, especially at the level
>   talk about food, furniture, animals & other things that are
>   central to the experience of young children.
>

I'm suggesting that by looking at the hierarchy of perceptions, you can find some constraints which the language structure then doesn't have to provide. If verbs, for example, refer to perceptions at the event and transition level, then there's a way to identify a verb without formal reference to its relation to other perceptions (SUBJ and OBJ). Language structure can't violate perceptual structure (if meaning is to be signified), so you can't have just any old rule.

Greg and Pat Williams (920219) --

>Question: Where does the 5% (or better) accuracy figure for PCT models
>come from? Which models, what data? I (Greg) admit to surprise regarding
>this figure for the Little Man, as I didn't know Bill had compared it to
>"real" data.

This figure refers to models of visual-motor tracking performance, where we have accurate comparisons of model and human data. I have done these comparisons using many unpredictable disturbances, changes in handle sensitivity, nonlinear connections from handle to cursor, and multiple disturbances acting at different places on the display. Tom Bourbon has done the same things, plus matching models of multiple control systems (in one person and in different people) in cooperative tracking tasks. Rick Marken has found similar results in two-dimensional tracking tasks,including tasks with unexpected reversals of handle effects on the cursor.

In all these cases, overlaying the simulated behavior on traces of the real behavior yields deviations that are nowhere more than about 5% of the range of movements observed. These models are truly predictive, in that the parameters measured using a simple preliminary experiment are then used unchanged to predict behavior under new conditions, with new patterns of disturbance and even new disturbances not present in the original measurement. The figure of 5% is a conservative one for the accuracy of prediction.

The Little Man hasn't been compared with quantitative human data (by me). It does predict some aspects of pointing behavior quite accurately: the finger follows a slowly-moving target just as a real person's would -- the tracking error in both cases is essentially zero, and is certainly less than 5 per cent of the range of target movement in three dimensions. An open-loop model of pointing behavior published in Science a few years ago achieved a minimum pointing error of something like 4 degrees in direction and 4 centimeters in position, if I remember correctly -- with a STATIONARY target.

The dynamics of the Little Man are not quite right: when extremely fast movements from one place to another occur (step-changes in reference signal), the trajectory of the fingertip is not an accurate reproduction of real movements. By adjustment of loop gains, the response of an outstretched arm to switching gravity on and off can be made pretty realistic, but I have no quantitative data on that. The behavior of the fingertip near the target after a jump in target position is qualitatively very realistic -- the slowing down, the little wobbles, and finally the touching. The fingertip does not avoid passing through the target, because there are no touch sensors in the arm. To compare this aspect of movement with real behavior, you'd have to use a target superimposed on the behavior space using a half-silvered mirror.

I believe that if you read what I wrote more carefully, you'll see that I didn't claim that every control-system model I've proposed matches human behavior within 5 per cent. This was intended to apply only to those simple models we have been able to test quantitatively against human behavioral data. If I didn't say that, I certainly should have. I usually remember to include that limitation.

I've looked over the papers you sent, for the "Unpublished papers" volume. I hate them all (except maybe the "New Epistemologists" one that will make sure I'm never invited to another cybernetics conference). This results from being very bored with reading my own writing. The worst one is "Using control theory in higher realms of psychology," which rambles and rambles to no good end. Trash-can it. The paper on transfer functions refers to data taken 4 computers ago that I can't find and wouldn't want to recreate. I'll write this one again some day; it isn't worth printing as is.

As to the rest, use your own judgment or yours and Tom's. I hope you'll put the oldest stuff last, because if it has any interest at all it certainly isn't in the writing, and the concepts are pretty dated. I particularly hate the papers from the days when I underlined half of the words. The frontispiece is OK. I wrote that bit for Dag Forssell to use, so it sounds pretty promotional. But I believe it.
When it's all printed and bound I will probably like it better. With my present attitude I'm the wrong person to ask what is worth preserving.

Best to all,

Bill P.

Date: Thu Feb 20, 1992 2:32 pm PST
Subject: More segue

From Greg Williams (920220)

In considering Rick's request for ideas on what sorts of models might garner attention for PCT, I came up with the following. I think that various audiences (neuroscientists ("NS"), psychologists ("P"), roboticists/AI researchers ("RAI"), and sociologists -- including economists, political scientists, and management scientists in the last group ("S")) would be most impressed by matchings between particular sorts of data and model dynamics. In the hope that Rick needn't reinvent wheels already extant in PCT models, I have indicated my impressions of how well the existing PCT models seem to me to fit the various kinds of data. A caveat: My evaluation of the data fits for existing PCT models isn't up-to-date, because I don't know what data they have been tested against recently, and I limit my evaluation to three PCT models with which I am most familiar: single-person tracking ("Tr"), the "little man" ("Arm"), and gatherings ("Ga").

1. Neurophysiological data. I don't really want to reopen this can of worms. I'll just note that this is a type of data, and that I think audience interest in a model's conformity to such data would be highest for NS, less for P, and least for S; I suspect some RAI would be interested, others wouldn't. Suffice it to say that opinions among PCTers regarding the actual and potential fit of Tr, Arm, and Ga to neurophysiological data are divided.

2. Behavioral data. By "behavior," I mean the PCT sense of the word: achieving invariable ends via variable actions. A model fitting such data well would reliably achieve certain goals in the face of disturbances. All three of the PCT models, I think, do this, and it is difficult to fault them. One technical point which could indicate potential (minor) problems with Arm (and Ga? I'm not sure) is that because of the particular control mechanism (non-integral) used, there is a slight steady-state error -- the goal is not quite reached; this error can be reduced by raising the gain, but stability difficulties might result; regardless, the error problem can be solved by slight changes in the model. NS, I think, would be interested in this fit only to a limited degree, P would be interested to a considerable degree (more so as their specialty was less "physiological"), S would be interested to a great degree, and RAI would also be interested to a great degree. However, I think that S and RAI would tend to split on issues of instantiating the models: RAI would be more interested in models which behaved in the real world (robots) than in purely computer-based simulations; S would be happy with computer-based simulations which behave.

3. Action data. By "actions," I mean the PCT sense of the word. Such data used to constrain models would be of the following form: specific, detailed spatiotemporal trajectory plots recorded during behavioral sequences which were subject to the same disturbances. I suppose there could be an issue over how to make sure the disturbances are the same in successive data-gathering "runs," but I'll bypass it by restricting my notion of "action data" to data sets with little variability (think, i.e., of the pro baseball pitcher's windup and delivery of his "trademark" fast ball, into which the experimenter introduces, each time, the "same" disturbance of a jerk on the ball of a given force at a given "place" in the windup). With respect to this sort of data, to my knowledge, Arm is untested (well, actually, I attempted to match planar trajectory curves of real arms -- as recorded by MIT researchers -- to

trajectory curves of my (DESIRE) implementation of Arm, with negative results) and Ga is untested, but Tr matches the data quite closely. Unfortunately, I believe that the close match of action data by Tr (as known to me -- maybe this is no longer true) is due to the easiness of tracking in the experiments used to generate the data. The subjects closely tracked the cursor, as requested, and any (control) model which did the same would have high correlations with the data; the proportional control model used has higher correlations than some of the other kinds of (control) models, but it has lower correlations than others; what is needed to distinguish among the possible control models is data which is more sensitive to model type. One way to generate such data is to use more difficult tracking conditions. The quite complex "human operator" tracking models have been narrowed down by reference to such data. I think NS would show considerable interest in models which fit action data -- but probably only if the data used allowed narrowing among control models. Ditto for P. S probably wouldn't care much. And neither would RAI.

So, Rick, you see that my advice to you about "interesting" PCT models must reflect the prior question "Interesting to whom?" Here are some possibilities which build on the existing models:

1. Enhance tracking model to cover "difficult" conditions (and maybe even the effects of practice/familiarization). The "human operator" models are partly curve-fit, and an entirely generative PCT model would probably be received warmly by NS and P, though RAI and S probably wouldn't be much interested.

2. S are probably already satisfied with Ga. Similar models could be made for sub-disciplines (economic exchange models, etc.). Some (so far, unfortunately, only a few) P would be interested. Robot implementations (and to an extent, more "realistic" computer simulations -- with more disturbances) would interest RAI (and some P). NS probably wouldn't be much interested.

3. Test Arm (if this hasn't been done already) and Ga for fit to action data. If the fit is good, NS and many P would be interested. If it isn't, work on the models to improve the fit. This is open-ended -- the more stringent the action constraints, the more interest, probably. Robot implementations of Arm are difficult (the perceptual problem) and costly, but ultimately are the way to generate much interest from RAI.

These options are not meant to be constraining in any way. Your "flashy" (no cut intended) ideas about a juggling simulation, etc., have appeal. I think you could see how they would fit in the typology (perhaps modified as you see fit). My guiding principle was to suggest ways to use whatever you can of the already existing PCT models.

Constructively (I hope),

Greg

Date: Thu Feb 20, 1992 2:46 pm PST
Subject: PCT as dogma

[From Rick Marken (920220b)]

Alan Scrivner (920220) says:

>   Does PCT provide a theoretical

>basis from which one can do better than guessing when the data runs
>out? Will it allow a basis from which to launch a search of the space
>of possible control configurations for the development of intelligent
>agents which exhibit particular behaviors?.

What do you mean by "behaviors"? If you mean "controlled consequences of action" then the answer to your question is a hearty YES.

>   Not if open-minded, scientific
>investigation is replaced with dogmatic attacks of the form featured here.

Well, I don't know that they were "dogmatic attacks" or even "attacks". My main criticism of Beer (to be endlessly redundant) is that he is not oriented toward modelling control -- though Greg and Pat argue that he is. I agree that his bug controls some variables. I just don't see the control exhibited by the bug as being based on an understanding of basic principles of control. I think all of the "attackers" acknowledged and celebrated the engineering achievement represented by Beer's bugs. They just didn't see it as representing any fundemental push forward in our understanding of the organization of living systems. But I think all of the "attackers" -- really critics -- are willing to change their minds if convinced. But we can have an opinion, can't we?

Dogma means unwillingness to change one's beliefs based on evidence. I'm willing to be changed. I'm not convinced that Beer has done anything fundemental as far as modelling the kind of behavior I am interested in -- purposeful behavior -- but I'm willing to be swayed. I just havn't been swayed yet. It's not simple to change someone's mind -- PCTers are sure aware of that. I've seen PCT models do a lot of powerful things and they impress me as being deeper than the Beer model. It seems that the main argument in Beer's favor is his fidelity to neurophysiological "truth". I'm not convinced that the Beer model is any more physiologically "truth-ful" than the PCT model. I've read his book -- I think I've made a good faith effort to try to understand what he is doing. I disagree with his fundemental assumptions (at least as they are expressed verbally) about how behavior works. But the bug does do some controlling -- and I understand how it does it -- to some extent.

I do feel bad if it seems like we are attacking Beer; it seemed to me like we were just criticising -- isn't that one role of science? I certainly don't dislike Beer's work; I'm just not impressed by it. I think you might want to take a look at Bill Powers' "Little Man" to see an example of modelling that I do find impressive. You might disagree -- vibrantly -- but I won't consider it an "attack" -- just an opinion.

By the way, you might also try my hierarchical control spreadsheet model; it does some pretty intereting things, though there are no graphics (the behavior is just numbers). But once you understand what it's doing you might see some of the potential power of Powers' model.

Best Regards

Rick

Date: Thu Feb 20, 1992 3:27 pm PST
Subject: models

[From Rick Marken (920220c)]

Bill Powers (920220.0900) very clearly and concisely points out the fundemental difference between PCT and non-PCT approaches to modeling behavior. I've been trying to make this point all along -- with all kinds of different word combinations (attn: linguists) but Bill's nice concrete description of how to build a PCT model makes my point perfectly:

>In defining the behavior of models in simulated environments, it's
>important not to use the modeler's perceptions in specifying what the
>system is to "do." The model would not "search for food" -- that is an
>observer's impression. The model would try to bring some sensory indication
>of the presence of food to a reference level. Its actions might look to a
>human observer like "searching" and the object might appear to be "food,"
>but unless the model itself contains such notions, this can't be a correct
>description from the model's point of view. Even "edge following" is a
>human interpretation -- the cockroach doesn't perceive anything
>corresponding to "edgeness" or "following." It just controls its antenna
>touch signal. By doing so it appears to us to "follow an edge," "navigate
>around an obstacle," and so on. That is how a human onlooker interprets the
>results, anthropomorphizing.

>What we're trying to do, I think, is to construct a system that controls
>for various perceptual variables by use of its actuators. In a given
>environment, this will lead to much behavior that the human observer can
>see. The human being will see all sorts of high-level variables in the
>movements of the model, and marvel at the "intelligence" demonstrated by
>it. This happens all the time with the Gatherings program.

>   The point of the control-system model is
>to show that you can get a lot of apparent complexity of behavior out of
>controlling just a few simple perceptions in a complex environment.

I'm sorry to quote so much -- but I can't think of a better way to express the difference between the PCT and the Beer approach. Again, not that Beer's work isn't a great achievment. But, can't you now see the rather profound difference between the attempt to model "behavior" and the attempt to model the "control of perception"??

Best regards

Rick

Date: Thu Feb 20, 1992 6:11 pm PST
Subject: Beer debate

[From Rick Marken (920220d)]

Well, I'm going on jury duty next week so what they hey:

Dag Forssell (920220) takes me to task for misrepresenting the status of the Beer bash (er, debate):

>Greg and Pat have consistently argued, most clearly in their post (Greg
>& Pat 920218), not that

>but that Beer deliberately and explicitly models CONTROL.

Explicitly, perhaps, (in that the bugs do control some variables -- like

the presence of an "on" signal for leg extension) but not deliberately.
I could be wrong -- just test for Beer's controlled variables. I bet
he is trying to make the bugs appear to be doing certain things -- not
to be controlling particular perceptual variables -- see my previous post.

> Beer's sin is
>not to use PCT terminology.

No sin -- Beer is just making a mistake. He seems to think behavior is what the observer of the the organism sees. He gives no evidence that he understands that what an observer sees as "behavior" is just a perspective on the organism's efforts to bring it's own perceptions into correspondence with its references for those perceptions. I think this misconception is a lot deeper than a terminological difference. Again, I might be wrong. If so, I'm happy to be persuaded.

I have nothing against Beer or Behaviorists or anyone else. I'm just stuck thinking that organism's control perceptions. I've had Behaviorists and AIists and whomever say, yeah, yeah that's right and then give every evidence that they had no idea what that meant. I can't honestly say "yeah, you're right" to Beer if he's wrong (or, at least, if I think so). I'm not mad at him, any more than I'm mad at people who think the sun goes around the earth (if there are any). I just think they are wrong and I think I can show them why (that's what all that research I've done is about -- it's obviously not done to get popular).

Best regards

Rick

Date: Thu Feb 20, 1992 10:44 pm PST
Subject: BEER stuff; model data

[From Bill Powers (920220.1800)]

Alan Scrivner (920220) --

>Does PCT provide a theoretical basis from which one can do better than
>guessing when the data runs out? Will it allow a basis from which to
>launch a search of the space of possible control configurations for the
>development of intelligent agents which exhibit particular behaviors?. Not
>if open-minded, scientific investigation is replaced with dogmatic attacks
>of the form featured here.

The elimination of dogmatic attacks is surely in all our interests. Not only do such unreasoned attacks create equally unreasoned defenses, they create a smokescreen that conceals real and important differences of approach, and even factual disagreements. Real disagreements call for intense investigation because somewhere behind them are contradictory assumptions about the nature of behavior.

There seems to be a direct conflict of assumptions between G&P Williams, and (to pick a scapegoat) Rick Marken concerning whether Beer is investigating "control". This is a confusing issue, because in at least two respects -- obstacle avoidance and food seeking -- the Beer _model_ does in fact control something. In the first instance it controls the signal from an antenna touch sensor, and in the second, the signal from a smell sensor (in fact it controls the latter in two ways: total smell and left-right difference). The question Marken is raising is whether Beer recognizes these control processes and has deliberately constructed his model so as to carry them out, or whether he has some other conception of behavior and in the course of implementing it happened to produce some control systems by accident.

Greg and Pat, on the other hand, assert that the Beer model actually achieves control but does it in a way different from the way outlined in the basic or "canonical" CT model (which consists of a perceptual function, a comparator, and an output function). This is basically a disagreement about what the term "control" is to mean.

The walking pattern is an example. Here we must ask whether the legs and the circuit that runs them go through a controlled disturbance-resistant pattern, or whether they act simply like a complex wheel run by a motor, the speed changing but not itself being controlled. Beer himself suggests the latter: when a physical model using this circuit is pushed (with higher circuitry disabled), it goes through the same walking sequence as when neural commands operate it, as the limit-sensors switch the legs from forward to reverse motion. There does not seem to be any resistance to the push capable of cancelling its effects, or any attempt by the system to enforce a stationary state. As the essence of control is resistance to disturbance (usually very pronounced), this system does not seem to meet the requirements for being a control system or accomplishing control. It is probably better described as a complex output function and effector. That doesn't make it any less interesting.

There are really no fundamental disagreements between CTers and Beer, save for the conception of what is to be explained by a model. All CT modelers, as far as I know, assume that the functions in living control systems are carried out by neural networks. I've always assumed this, as BCP shows, even though I haven't tried to guess what the networks actually are. It's good that people are exploring at that level, and coming up with plausible if not firm answers. The way we will settle the disagreements with specific models outside the CT framework is through behavioral tests (assuming that firm neural models are some distance in the future). While we can't converge to a single right model with such tests, we can at least eliminate models that misbehave.

Greg Williams (920220) --

Good summary of types of data.

>1. Neurophysiological data.

The CT model is consistent with neurophysiological data and models. A comparator, for example, is a Beer neuron with one positive input and one negative input. A composite comparator, capable of producing error signals that pass through zero, is made of two Beer neurons with the positive and negative inputs interchanged, and the error signal consisting of the pair of output firing rates. A simple output function is a Beer neuron with amplification and a slow integration. The only missing function is the perceptual function, which save for trivially simple cases (intensity perception, perception of weighted sums) is beyond our ability to model because we don't know how it works for higher systems. Of course for more complex functions we have to use more Beer neurons -- but for any reasonably simple control system such as the arm model, there's no linear function we can describe that couldn't be modelled with Beer neurons. We could even make a rate-of-change sensor by using a Beer integrator in a negative feedback loop. The Arm model could be implemented completely with Beer neurons (save for the environment and the details of vision).

2: Behavioral data:

>... because of the particular control mechanism (non-integral) used, there
>is a slight steady-state error -- the goal is not quite reached; this
>error can be reduced by raising the gain, but stability difficulties might
>result; regardless, the error problem can be solved by slight changes in
>the model.

The arm model uses integral output at the visual level -- error is ultimately zero. It isn't necessary for the kinesthetic systems to achieve zero error, because higher systems, which integrate, just keep changing the reference signals until the result is right. Ditto for the Gathering model; the error signals are converted to an increment or decrement in the direction of travel, so the actual angle of travel is the integral of the error. No steady-state error.

Generally, there's really not much difference between a control system with a gain of 30 and one with an infinite or integral gain. With a gain of only 30, the error signal is 3 percent of the reference signal with a full-scale disturbance -- you wouldn't see the error in ordinary behavior.

>3. Action data. By "actions," I mean the PCT sense of the word. Such data
>used to constrain models would be of the following form: specific,
>detailed spatiotemporal trajectory plots recorded during behavioral
>sequences which were subject to the same disturbances.

It's hard to get meaningful action data. In the normal range of control, where disturbances are of the size and speed for which the system has become adapted, actions simply mirror disturbances with very little error. Only when you start making disturbances very large and very fast does the action pattern depart significantly from the disturbance pattern -- and then, of course, you aren't seeing good control any more. In the range of good control, sizeable differences in sensitivity and linearity in the output function are hidden by the control process; they make essentially no difference in the mirror relationship between action and disturbance, or in the way the controlled variable follows the reference signal. When you get outside the normal control range, particularly in regard to speed, the feedback no longer enforces the mirror relationship or the tracking of the reference signal, and you see the effects of variations and nonlinearity in the output function -- which can be different for every individual.

>Unfortunately, I believe that the close match of action data by Tr (as
>known to me -- maybe this is no longer true) is due to the easiness of
>tracking in the experiments used to generate the data. The subjects
>closely tracked the cursor, as requested, and any (control) model which
>did the same would have high correlations with the data; the proportional
>control model used has higher correlations than some of the other kinds of
>(control) models, but it has lower correlations than others; what is
>needed to distinguish among the possible control models is data which is
>more sensitive to model type.

The second-to-last demo in DEMO2 uses an adjustable integration factor with an adjustable leakiness to match model behavior to real behavior. It is definitely possible to make the model track considerably better than the real person. The RMS error between model handle and real handle is GREATLY reduced by adjusting model parameters to make it track less well. The last

demo adds a slowing factor in the perceptual function, and adjusting this
leads to a further reduction in RMS prediction error, with a comcomitant
increase in the errors made by the model. This last effect is not large,
but it is very consistent from one experimental data set to another.

Also, predictions of both handle and cursor behavior improve when going
from very easy (slow) disturbances to harder disturbances. The
systematic errors increase in relation to small random errors,
and the model can reproduce them better, with a sharper selection of the
right parameters. VERY hard disturbances reverse this improvement --
subjects start to lose control and most of them give up. The model never
gives up or reorganizes.

> One way to generate such data is to use more difficult tracking
>conditions. The quite complex "human operator" tracking models have been
>narrowed down by reference to such data. I think NS would show
>considerable interest in models which fit action data -- but probably only
>if the data used allowed narrowing among control models.

Demo2 allows four levels of difficulty of disturbance in the final
sections. All the effects you're talking about are shown, I think. Try it.

The real difficulty with action data comes in when you consider more
complex behaviors instead of simple situations with only one allowed
action. A control-system model could predict very well that you will drive
from home to work and end up in the right parking place. But the trajectory
you use to get there wouldn't match a model's trajectory unless the model
happened to choose exactly the same alternative behaviors that you did.
Even if you just had two control handles contributing to a single cursor
position, one in each hand (or two dimensions of mouse movement), the
relative movements of the two handles could be any combination at all. The
control model could predict the sum of the handle positions, but there
wouldn't be any basis for predicting how much each handle would contribute.

You can do some learning experiments with Demo1. Using the "Controlling
different variables" section, you can let a naive subject control one of
the variables again and again in a series of 30-second runs. After each
run, the program shows the correlations between handle and cursor, and
between handle and disturbance. The first correlation should gradually get
lower, and the second higher. A better experiment would also compute the
best model for each run -- I haven't set that up, but it wouldn't be hard.

Dag Forssell (920220) --

Thanks for upbeat remarks -- they help us remember our principles.

Best to all,

Bill P.

Date: Fri Feb 21, 1992 8:03 am PST
Subject: Clarification

From Pat Williams (920221)

>Bill Powers (920220.1800)

>Greg and Pat, on the other hand, assert that the Beer model actually

>achieves control but does it in a way different from the way outlined in
>the basic or "canonical" CT model (which consists of a perceptual function,
>a comparator, and an output function). This is basically a disagreement
>about what the term "control" is to mean.

I guess we weren't very clear about the differences we see between PCT modeling and Beer's modeling. We think that both use perceptual functions, comparators, and output functions. The main difference we see is in Bill's detailed hierarchical model, the higher level errors set lower level reference signals, while in Beer's bug, certain lower level systems are inhibited to change modes.

Best wishes,

Pat

Date: Fri Feb 21, 1992 8:20 am PST
Subject: Bridge construction

From Greg Williams (920221)

I think we're making some progress in steering toward more constructive topics.

Over the past couple of days I've had one of those aha experiences that seemed to reorganize a bunch of my notions, particularly as they relate to PCTers' relationship with non-PCT psychologists. What I realized, much more deeply than ever before, is that the most prominent PCTers tend to value models of behavior (the PCT meaning) and to dismiss models of action (again, the PCT meaning), while nonPCTers tend to value models of action and to dismiss models of behavior. Right now, I am convinced that this distinction goes to the heart of the less than enthusiastic enthusiasm for PCT among many psychologists, and why that situation isn't changing. This, of course, is an empirical hypothesis, to be tested, I suppose, using data from peer reviews and so forth. Right now, I simply put it forth as a plausible insight about what is going on.

In the past, I more or less glossed over Bill's claim that organisms generally maintain control pretty well -- their error signals are never very big. I did find it somewhat hard to believe, because, for instance, when I want a drink of water, I do have to go through a period of transient action before that want is satisfied (the water doesn't just gush into my throat when I want it to -- I have to go get the water, which takes some time). To a large degree, I suppose (now) that I took Bill's claim as a bit semantically confusing (since there are obvious transients in organismic control), but not really of huge consequence. I could live with it, because it seemed to be a pretty innocuous assumption. No longer do I consider it innocuous. I now think it is the basis for what divides PCT from non-PCT in practice. Please note, Bill, that the worth or correctness or utility of your claim isn't of concern here -- if the belief in it is present and it guides certain assumptions made by PCTers, then that's enough to support my hypothesis. And I'm not trying to say that the nonPCTers AREN'T deluded, in various ways. This is an attempt to understand the basis for the situation, not to prescribe possible remedies.

First, consider non-PCT psychologists. Many of them, I believe, do indeed think (in an ill-defined way) that organisms control -- that there are what amount to reference signals inside, and that there are negative feedback loops. This, in a sense, is what the "cognitive revolution" has wrought.

Others (generically, behaviorists) don't want to think at all about such models, because they think they don't need to (they are simply wrong if they want to get very far). Good luck convincing this second group. A third group aren't prebiased, and some of them have ended up as PCTers. But the first group, I believe now (ta da!), think that the fact of control is essentially trivial, well-known, generally accepted, AND NOT WELL-UNDERSTOOD. What they are looking for which will seem non-trivial to them are detailed models for how control occurs -- the mechanisms involved in very particular bits of control behavior. And the way they think that such mechanisms can be figured out is by making models which show "correct" (empirically supported) actions during those very particular bits of control behavior. Behind this belief are, for example, notions in control engineering that the structure of a dynamic system can be identified by examining its transient (not steady-state) behavior. (Yes, Bill, I am glossing over technical details here, but I don't think my characterization is ridiculous.) So they think that models which predict actions under precisely controlled circumstances are great stuff, and models which control (disregarding the transient details) are ho-hum.

Prominent PCT model builders have the reverse priorities: models which control (in general) are important, and their specific transient actions aren't important to worry over (for a variety of reasons, the basic one being Bill's claim that transients aren't very important in organismic control; another is the idea that prediction of actions is difficult for various reasons). So they build models which control, and usually don't get into the details of particular actions matched by those models.

The result, I think, is that the models put forth by PCTers generally are not seen as very important by many non-PCT psychologists, and PCTers fail to see why this should be so.

Some support for this "explanation of The impasse" is found in the interest generated among a number of psychologists by PCT tracking models which predict actions quite accurately.

As Warren McCulloch was wont to say, "Don't bite my finger, look where I'm pointing." (And I'm not pointing the finger at YOU! Honest!!)

Greg

Date: Fri Feb 21, 1992 8:26 am PST
Subject: LFG, PFG, PSG

[From: Bruce Nevin (Fri 920121 07:36:42)]

LFG (Lexical-Functional Grammar) appears to me to retain the problems of phrase-structure grammar (PSG) rewrite rules, and what it adds does not adequately address those problems.

First, the PSG problems. Rewrite rules do not classify words, they classify phrases. Thus:

```
  S  ->  (XP) NP VP
```

Any operator morpheme that can appear in a VP can also appear as an argument in an NP, accompanied by an indeterminately large number of other words in the VP and in the NP. These are specified by recursive rewrite rules such as:

```
NP -> A NP
NP -> N
N  -> book, dog, . . .
A  -> red, big, . . .
VP -> D VP
VP -> V
V  -> chew, read, . . .
D  -> quickly, late, . . .
```

These rules specify an indefinitely ramified branching hierarchy of phrase classes, with S (sentence) at the root and individual words at the leaves.

```
                        S
                      /   \
                    /       \
                  NP         VP
                 /  \       /  \
               A    NP     VP   D
               |   /  \   /  \  |
            small A   NP VP NP late
                  |  / \  | / \
                 red. . . . . . .
                    . . . . . . .
                    . . . . . . .
                  ..dog..chew..book..
```

But the elements above the leaves (S, NP, VP, etc.) have no perceptual connection with individual words (dog, chew, book) in themselves. The same individual operator morpheme may appear in VP or in NP ("the dog's chewing of the book angered her"), or in a sub-phrase under either of these (a modifier: "the chewing damage cost the most"). There can be no connection with nonverbal perceptions or semantics until "lexical insertion," that is, until after the abstract sentence structure is generated down to preterminals (the N, V, A, etc. to which the leaves of words are attached).

This necessarily reduces PSG-based syntax to a filtering function ratifying an unstructured collocation of words motivated by nonverbal perceptions. Just throw out words in whatever order they occur to you. If that order doesn't match a possible syntactic structure--you have to have ready access to the infinite set of possible syntactic structures, of course, indexed by leaf and not from the root down--if the perceptual order doesn't match any available structure, then reorganize. Reiterate random reorderings until you find a match. Possible, I suppose, but not plausible. Not when there are simpler alternatives available. And not when computer implementations founder on precisely this issue of selection restrictions.

For what is the basis for specifying phrase types? Selection restrictions. Limitations as to which words can go together with which other words. But selection restrictions are not firm, well-defined things. Whatever subclasses of words you define--human, animate, abstract, concrete, etc.--as soon as you say only +animate nouns can cooccur with a +animate verb, someone will think of a perceptual situation, in describing which you would use a -animate noun with a +animate verb. The teacup danced off the table.

Why? Because these features are not properties of the words. Nor are they properties of things to which words refer. They are byproducts of perceptual control by language users. And if a person wants to perceive a situation as if a teacup is animate and can dance off the table, that person will do just that, and no amount of expostulation about semantic features and selection restrictions will have the slightest effect. Selection restrictions and semantic features are linguists' generalizations, not internalized rules of grammar (universal or otherwise) that people follow in producing sentences. This is why Harris abandoned the search for semantic features or components long ago. (Though when I proposed doing componential analysis of dictionary definitions, he only said "others have tried and have failed, but you are welcome to try if you want.") And this is why "lexical insertion" and selection restrictions and the "semantic component" have been for 40 years something that generative linguists have always been going to get to real soon now. They still haven't a clue.

And these problems arise necessarily from adherance to PSG rewrite rules and the, phrase types, labelled bracketings and tree structures that they define. These tree structures have become so familiar to linguists that they think of them almost as properties of language. But they are not properties of language, they are only a particular descriptive mechanism, a mechanism that is not necessary--there are alternatives--and that introduces structure (restrictions) of its own not found in other descriptive mechanisms. Because its restrictions are not found in e.g. operator-grammar dependencies, these restrictions are artifacts of PSG and not properties of language (insofar as both descriptive mechanisms can account for the same utterances).

LFG introduces a refinement of PSG. To each phrase marker it adds an index specifying its syntactic function: SUBJ, OBJ, etc. Since it is the function of each phrase that is specified in this grammar formalism, it could perhaps more accurately be called phrase-functional grammar. But what does it mean that the first NP is subject of the VP? For example, in:

> The dog, apparently incensed at seeing the cup dance off the table, continued to attack the shards scattered in ignoble retreat across the floor.

At the level of NP and VP you don't know what words are involved. You don't know that the NP is

> The dog, [which was] apparently incensed at seeing the cup dance off the table

Nor do you know that the VP is the moiety of the sentence:

> continued to attack the shards scattered in ignoble retreat across the floor

Indeed, if you could know this (within the PSG system) it still would not be useful information, because you don't know that the relevant words for selection restrictions are "dog" and "attack". These are the head words in the NP and the VP, respectively.

The basic problem is that these rewrite-rule systems obscure the relations between individual words that must be brought into correlation with perceptions. Attempts have been made to overcome this fundamental deficiency of PSG. Socalled X-bar notation, originated by Harris in 1946 ("From Morpheme to Utterance" in one of the issues of _Language_ for that year) and taken up by Chomsky 20 years later, has given rise to head-driven PSG and its relatives. It was in recognition of the inadequacy of X-bar notation really to address the problem that Harris invented string grammar (center-and-adjunct grammar). And it was because of weaknesses of string grammar (it doesn't handle embedding well) that Joshi created his very successful Tree-adjoining grammars (TAGs), in which extremely restricted rewrite rules are combined with the adjoining rules of center-and-adjunct grammar. And for the same reasons, Harris abandoned string grammar and moved thence to the word dependencies of operator grammar. The relations of adjoining in the former and dependency in the latter are relations among pairs of words.

To address the same word-word : perception-perception association in a PSG-based system like LFG requires an inordinately complex superstructure of Principles and Constraints and other Generalizations. Much of the hoopla in the literature takes this form:

> The Friblewhisket Principle (FP) accounts for the fact that you can say sentences 3, 5, 7, and 13, but you can't say sentences 2, 4, 6, and 11. The FP is independently motivated (Chomsky (1987b), Newmeyer (1990)), so that is prima facie evidence that it is a property of Universal Grammar. That indicates that this explanation of the unacceptability of 2, 4, 6, and 11 is indeed the correct one.

If you look up Newmeyer (1990) you find another reference to Chomsky (1987b) and no new data. If you look up Chomsky (1987b) you find on p. 342 of the journal an avowedly informal and "suggestive" assemblage of familiar examples in a new combination, with a rough and "intuitive" formulation of the FP, with a promise to demonstrate it more adequately further on, but look as you might that promise does not seem to be fulfilled. (I think Postal called this the "Phantom Principle Move" in his hilarious "Advances in Linguistic Rhetoric." He gives examples of this and other vacuities people in the field have passed off as substance.) But there it is, reified: the FP, a sexy concept to be conjured with by writers of dissertations and conference papers--until it ceases to be fashionable and no longer demonstrates that one is in the know.

All of this could be written off as sour grapes, except for this: when you look at the same data from the perspective of an alternative theory, there is no problem, or the issue that seemed to be a problem is explained by something entirely different from the Friblewhisket Principle. This alternative solution cannot be stated in PSG terms without recourse to something like the FP, or its successor or predecessor, because something extra is always needed to make up for the inability of PSG tree structures to show dependencies among noncontiguous words.

But if you try to articulate an alternative perspective to "mainstream" linguists, they dismiss what you are saying as not interesting because it doesn't connect to the FP or the OGP or the ETC and so on. I heard a student of Wally Chafe say to an assemblage of linguists that the problem area addressed in a series of papers heard that day simply did not arise as problems in Chafe's "semantically based" approach. Similarly for adherants of Sid Lamb's stratificational grammar. So it doesn't just come up for Harrisians. You'd think that failure of phenomenon X to be visible from an alternative theoretical perspective might be a clue that X is an artifact of one's own perspective, a scratch on the sunglasses. But no, it is simply incomprehensible, or at least news of invisibility and not being problematic is met with uncomprehending silence. I think this is because the way to advance is to demonstrate that one knows, and can teach to students if one is hired, the latest and greatest of mainstream linguistics. Sick transit.

That is why I am writing a very constrained dissertation on phonology, which I believe I can do creditedly without feeling like a whore. The things that phonologists are worried about concern the control of perceptions constituting words and morphemes. Phonological issues are almost entirely below the level at which linguistic perceptions are correlated with nonverbal perceptions. I just won't publish anything about syntax and semantics until after the degree.

Now, it may be that

>f-structures include something very much like the operator-argument
>structure that Bruce Nevin has occasionally talked about

but why take on the added complexity? There are other ways of handling operator-argument dependencies that are much more direct and computationally simpler.

Sure, you could express the argument requirement of each word with
>. . . a `PRED' feature, that specifies a name of
>a concept, and the grammatical relations of its arguments,
but rather than the abstractness of

> PRED `Bite(SUBJ,OBJ)'

you might say

   bite(N,N)

That is, "bite" requires two words in its argument, neither of which requires any argument (its in the Onn class of operators). Not just any SUBJ or OBJ will do. (I have previously discussed how to handle apparent extensions beyond these dependency classes, in ways that are difficult for PSG-based systems to accomodate. I'll get into that again if you think I am hoist by my own petard re selection restrictions.)

>One can think of this as a scheme for supplying operators with their
>arguments, which is supposed to function under various perturbations
>of overt linguistic form that linguists think they know a fair
>amount about.

I think it is great that you have formulated a broad range of linguistic generalizations in a computer implementation of a small LFG system. I hope I can get time to look at it, and maybe even make suggestions how it might be modified. You might find the modifications less than optimal for a pedagogical tool, however. Students want to learn the standard, mainstream stuff, and feel cheated if they are given something else, no matter whether it works better or not. At the level of discourse prevalent in the field, no one will have any basis or criteria for noticing that it works better. Not yet, anyway.

        Bruce

bn@bbn.com

Date: Fri Feb 21, 1992 8:54 am PST
Subject: Re: LCS II

From Greg Williams (920221-2)

Questions for Bill Powers:

1. Are you writing a Preface?

2. Are you sending disk files of "Standing at the Crossroads" and "The X Phenomenon"? If you haven't already, don't -- but let me know so I can get them typed in, as I want to include them.

3. Which paper do you mean by "New Epistemologists"? I'll make sure it gets included (unless you were being ironic).

4. How about editing out reference to the diagram in "The Arm"? I don't think much would be lost.

I'll have some more questions before long. Please answer these at your earliest convenience. Thanks!

Best, Greg

Date: Fri Feb 21, 1992 9:01 am PST
Subject: Discourse

[From Alan E. Scrivner (920221)]

Is Beer's work a breakthrough?

To this I would have to say no. What I see he has done, is to take Brooks "subsumption" ideas and craft the lower levels in a neurophysiologically engaging way.

What did this buy him?

Plausability!!! He worked from the bottom up. Generating interesting (from the observers standpoint) behaviors from the design and interaction of simplified, although biologically inspired processing units.

It's just a cute graphical trick his bug crawling around the screen like that.

Darn right it is! Salesmanship in science, don't leave home without it. Look at all the money they threw in Brooks' direction because he made slick little critters that crawled around the floor. One of his students got a PhD designing a wee beastie that was supposed to wander the lab picking up soda cans. I have heard that it actually got one, but that was by mistake. Beer is getting the attention he is in a large measure because of his engaging graphics, not the underlying mechanics.

How far can Beer take his technique?

Doesn't look like he'll go much further. I heard some time back that he is using genetic algorithms (GA) to evolve better leg controllers. Whenever one has to result to design by random mutation, my guess is that he's gone about as far as he can go. He has no theory to use as a design base so he uses "magic". This doesn't seem too far afield, since his origional purpose was to do what he has to do to get the behaviors he wants to observe.

So how could you do better?

Well, for starters take Beer's approach: Design in a neurophysiologically pleasing way.

Why?
Plausability of design from fundamental principles.

Will you use EXACTLY modelled neurons?

No. Every model makes assumptions and applies simplifications. But, the neurons should be more than the trivial threshold adders currently in use by the neural net crowd.

Why?

So they can be added to as more data on how real neurons function becomes available and because as with Beer's neurons modelled as differential equations, I believe the nonlinear effects are important.

What else?

Well, we have our basic units. Now we connect them into control centers (again neurophysiologically inspired) and build a hierarchy of such centers using PCT as a grounding theory to aid in the design.

What do we end up with? What Will it look like? and What will it do?

That depends on which of Greg's pigeonholes you call your own: (neuroscientists ("NS"), psychologists ("P"), roboticists/AI researchers ("RAI"), and sociologists -- including economists, political scientists, and management scientists in the last group ("S")). Personally, I fall into the RAI class and would like a simulator that could somehow be connected to a mobile robot. I think the feedback loop between realworld/virtualworld would be the most interesting and if PCT is of any value at all, my thoughts are that it would be of most value in a noisy environment. All suggestions welcome.

Date: Fri Feb 21, 1992 9:36 am PST
Subject: Bridge construction

[From Rick Marken (920221)]

Halleluja. I think Greg Williams (920221) has had a really remarkable insight
that must be taken into account by PCTers who want to be understood:

Greg says:

>Over the past couple of days I've had one of those aha experiences that seemed
>to reorganize a bunch of my notions, particularly as they relate to PCTers'
>relationship with non-PCT psychologists. What I realized, much more deeply
>than ever before, is that the most prominent PCTers tend to value models of
>behavior (the PCT meaning) and to dismiss models of action (again, the PCT
>meaning), while nonPCTers tend to value models of action and to dismiss models
>of behavior. Right now, I am convinced that this distinction goes to the heart
>of the less than enthusiastic enthusiasm for PCT among many psychologists, and
>why that situation isn't changing.

Yes indeed. Your observation, coupled with Bill's (that a lot of what is
called "behavior" by non-PCTers is really actions - PCT meaning -that may
or may not be involved in the control of perceptual variables) can be
parlayed into a way to make the PCT approach clearer to non-PCTers.

>First, consider non-PCT psychologists. Many of them, I believe, do indeed
>think (in an ill-defined way) that organisms control -- that there are what
>amount to reference signals inside, and that there are negative feedback
>loops. This, in a sense, is what the "cognitive revolution" has wrought.

Good point.

>   But the first
>group, I believe now (ta da!), think that the fact of control is essentially
>trivial, well-known, generally accepted, AND NOT WELL-UNDERSTOOD. What they
>are looking for which will seem non-trivial to them are detailed models for
>how control occurs -- the mechanisms involved in very particular bits of
>control behavior. And the way they think that such mechanisms can be figured
>out is by making models which show "correct" (empirically supported) actions
>during those very particular bits of control behavior.

Yes!!!

>   So they think that models which
>predict actions under precisely controlled circumstances are great stuff, and
>models which control (disregarding the transient details) are ho-hum.

Right on!!

>Prominent PCT model builders have the reverse priorities: models which control
>(in general) are important, and their specific transient actions aren't
>important to worry over

Yes!! I think this is an brilliant point. I think PCTers must show that,
while a system is controlling inputs it will be using what appear to be very
complex actions (Greg's "transients" -- which could be responses to transient
disturbances as well to transient error caused by sudden referecne changes--
as well as other side effects of the means used to control perceptions).
This will be the goal of my next piece of research -- show how appaerntly
complex "actions" (which non-PCTers call "behavior") can result from control
of simple variables.

>The result, I think, is that the models put forth by PCTers generally are not

>seen as very important by many non-PCT psychologists, and PCTers fail to see
>why this should be so.

You are absolutely right -- and I think I've now had your epiphany (thanks, Chuck T. for the word). I now understand (and empathize and even agree with) your attraction to the Beer model. Beer has designed a system that can be seen as exhibiting very complex "behavior" -- while what it is actually "doing" (ie controlling) is very simple -- it is controlling antenna signals and leg extension sensors. The Beer bug has very few and very simple controlled variables -- but an observer will see the bug as exhibiting all kinds of complex behaviors -- edge following, searching, feeding, etc. If one could make this fact about the behavior of the bug clear in the program write-up, it would make the Beer Bug a GREAT demo of PCT.

Thanks, Greg, for helping me understand the value of Beer's work and also for suggesting what I might do for my next research project.

Best regards

Rick.

9202D CSGnet

Date: Sat Feb 22, 1992 11:52 am PST
Subject: Re: psychologists and quantitative thinking

[Martin Taylor 920222 14:30]
(Dag Forssell 920220)
>
>A liberal arts education, including psychology (I don't know this, but
>it is my paradigm at the moment), does not teach or encourage
>quantitative thinking.
>

I suggest you change your "paradigm" (I assume you mean "understanding"). If there is a characteristic problem of experimental psychology, it is an overemphasis on the numeric and mathematical aspects of data acquisition, data analysis, and data interpretation, to the detriment of thinking about the qualitative aspects of what it all means. Of all the sciences, I suspect psychology has been the most diligent in developing mathematical methods of handling difficult data. My belief is that this is so because it is working near the boundary of where the degrees of freedom for observation are of the same order of magnitude as the degrees of freedom of the system being observed. In physics, the systems are so simple that a human mind can more or less comprehend them, and relatively simple mathematics can be used to make excellent predictions far from the actually observed variables. In psychology, many different observations at different times, under different environmental conditions, have to be amalgamated into some kind of abstraction before one can see any kind of consistency. This is true whether you are a classical S-R psychologist, a PCT psychologist, or any other kind. The success of a theoretical position can be measured only in the stability of its abstractions under the stress of environmental disturbance. For simple motor movements, PCT is quite successful in this way. Its principles may well lead to effective ways of combining observations so that the observer can reasonably say "these represent the same thing as those do," which is the toughest job in an experimental or observational science.

Psychology seems to me to be so dominated by the quantitative that the ability to step back and see the principles that organize observations is to be highly prized. That's why I honour Bill Powers' insights. He found a view that allows us potentially to go a long way further in predicting what we will observe in a living system.

You can fault psychology for not finding the grand qualitative principles, but not for failure in teaching or thinking quantitatively. (Of course, none of this applies to every individual psychologist, so I guess Bill wouldn't accept it as a statement of principle, or as useful description of psychology). A lot of kids go into psychology thinking that they can avoid mathematics. Most universities try to disabuse them of this idea quite quickly, if the kid intends to study psychology seriously.

(I guess I should qualify the last paragraph by saying that it applies only to the universities with which I am familiar, given the history of my claims about the extinct status of dinosaurs).

Martin

Date: Sat Feb 22, 1992 12:32 pm PST
Subject: BEERBUG design; bridges

[From Bill Powers (920222.1000)]

Pat Williams (920221) --

>The main difference we see is in Bill's detailed hierarchical model, the
>higher level errors set lower level reference signals, while in Beer's
>bug, certain lower level systems are inhibited to change modes.

It's true that the only communication between levels (downward) in the basic HCT model is through the reference signal. I've speculated about other modes such as gain adjustment (which in the limit is gating), and Rick and I in one model used gain reversal. But some of these effects can be achieved in the canonical model if we just consider some details of neurology.

It's possible, in a neural model of control systems, to turn off a control system just by setting its reference signal to zero. This is because a single neuron is a one-way device -- its output can't go negative. If a one-way comparator gets an excitatory reference signal and an inhibitory (negative current) perceptual signal, it will never output an error signal if the reference signal is zero, because with only an inhibitory input it won't generate an output no matter what the size of the perceptual signal. So a single Beer neuron used as a comparator can be turned off by setting its positive input to zero, provided the other input signal is inhibitory (negative). Note that "inhibition" has an ambiguous meaning: it could mean simply a negative input current, or it could mean gating off the output. I'm using it here in the sense of a quantitative negative input current.

When two-way action is required (as in limb movements), there can be two control systems, each using a one-way comparator, which produce positive output signals that have opposing effects (as in biceps and triceps). In order for control to pass through zero smoothly, the reference signals for both systems must be greater than zero. The two error signals create opposing muscle tensions. This is known as "muscle tone." The reference signals for these two systems would then have to work as a pair, one increasing while the other decreases. Any common-mode change would just raise and lower muscle tone. Of course with nonlinear muscles, the effect would also be to change differential muscle sensitivity, so adjusting muscle tone also affects the loop gain of the paired systems considered as a single system. It's therefore possible to turn systems off and to adjust the gain of paired systems just by varying reference signals.

In Beer's model, there is also control by changing reference signals. This is somewhat obscured because Beer thinks of the reference signals as gates that turn on the control systems, or else hides them inside his neurons. Consider "appetitive" behavior (diagram p. 129 -- I have the book now). There's an "energy sensor" that detects energy level, and when the energy level falls, "so does the activity of its energy sensor. This decreasing activity gradually releases the spontaneously active feeding arousal neuron from inhibition" (p.128). The feeding arousal neuron is therefore a comparator, the spontaneous level of firing being the energy level reference signal (set, no doubt, by a positive "intrinsic current") and the inhibitory perceptual signal varying smoothly with energy level. The output of the feeding arousal neuron is the error signal (one-way) which connects to the "search command" neuron, enabling the right and left turn neurons which are affected by perceptual signals from the left and right odor strength neurons. This turns on steering and brings the animal toward any food that's present.

Skipping the lower-level details, the result is to raise the energy level, bringing the perceived energy closer to the (fixed) reference energy. When the error at the highest level is zero again, the error signal goes to zero and directional tracking is turned off. So we have here the major high-level control system.

Among the skipped details is the fact that when the animal arrives at food, the presence of chemical signals enables the consummatory command, which causes food to become ingested. On p. 143, consummatory feeding is shown as inhibiting the appetitive system -- this is an ad-hoc design required by the lack of forward speed control to stop further search while the animal is at the food. The biting movements put food into the body, which is then metabolized (presumably) to raise the energy level and complete the highest-level control loop. The "overeating" phenomenon, by the way, can be modeled as a delay between ingestion of food and the subsequent rise in energy level. No "overeating neuron" would be needed.

Beer missed a bet by not using the sum of left and right odor signals. This sum represents total odor strength and thus distance from the food. The output of the feeding arousal neuron could set the reference signal for sensed odor strength, raising it when the animal is hungry (perceived energy level less than reference energy level). The difference between total actual odor strength and reference strength would be the error signal that varies the speed of forward walking. The same reference strength signal could enable the steering system also -- an on-off gating signal, if you want to retain that part of the model. There is, however, no need to disable the appetitive system to let the consummatory system work: as the animal approaches the food, it will slow down and stop when the odor strength reaches the reference level. Eating will then lower the arousal, which will lower the speed reference signal, and the animal will stop moving forward or move forward only slowly (a good thing if it isn't to keep biting at the same place over and over). The odor-based steering system can be left turned on all of the time.

There is really no need for a separate "wandering" system. The forward movement component of wandering results from raising the reference signal (or just the "command" if we consider the locomotive system to be an open-loop output function) for forward speed when sensed energy falls below the reference energy level. If the level of odor is very low, the signals from left and right odor sensors will be simply the noise-level in those receptors (which can be simulated), and directional turning, which can be activated all of the time, will be random. No ad-hoc "random burster neuron" is needed. So wandering will naturally take place when the animal is hungry and there is no nearby odor source. If the wandering brings it within odor range of food, the left-minus-right signal will rise above the noise level and systematic steering will commence. As a result, the total odor strength will begin to rise, although because of the inverse-square law it will rise rapidly only in the near vicinity of food: speed will remain nearly constant until the animal is close to the food.

Similarly, edge-following doesn't have to be disabled when the other systems are active. The total steering command can be the sum of error signals from the odor-direction sensing system and the touch-direction sensing system (the two directions of effect combined by subtraction or with left and right working separately). When there is no contact, steering toward food will work normally (or randomly if there's no odor). When there is contact, the touch sensors will cause veering away from the obstacle. Edge following will result when food is on the other side of the obstacle. When food is on the near side of the obstacle, the cockroach will not ignore it (if it's hungry) to go edge-following!

If you want a positive edge-following behavior to result, rather than just obstacle avoidance, you have to provide some reason for it -- an arbitrary turning-back toward the edge, ad hoc, is not a good reason, especially as you'd have to endow the cockroach with the ability to sense edges rather than just obstacles. We could ask what is better for a cockroach near an obstacle (like a baseboard) than away from it. One answer might be light level. There are light levels that decrease in the vicinity of an object, so the two eyes might detect (among other things) the difference in illumination between left and right visual fields. That difference, too, can contribute to the steering-direction signal (If you get the sign of the difference signal wrong, the bug will head toward higher light levels).

In the dark, cockroaches wander freely about the floor; in the light, they seek the shadows. When they encounter an obstacle before the light level has been reduced to or below the reference level for total light (left plus right), the collision-avoidance system simply makes them veer to avoid the collision, while continued seeking of a lower light level makes them veer back toward the obstacle. Thus they end up following edges, although they have no reference level or specific control system for following edges. Hunger, of course, also keeps them moving, although a similar system for total light-level control might also contribute to forward motion. In a dark enough environment, this hypothesis suggests, there will be no scurrying along the baseboards. Note that Chuck Tucker's cockroach-killing strategy takes advantage of the hypothesized seeking of a low light level.

So, Pat, I agree that Beer has not ALWAYS used control by reference level -- but I suggest that if he had, his model would work better and its circuitry would be simpler.

One omission from all Beer's diagrams is the external effect of motor actions on sensory inputs. He should draw arrows from activities like "biting" to the effect they have, here on energy level. The steering direction neurons affect the unbalance of odor signals. They also affect

the tactile signals. With these external connections explicit, the control loops become much more obvious.
-- Greg Williams (920221) --

>What I realized, much more deeply than ever before, is that the most
>prominent PCTers tend to value models of behavior (the PCT meaning) and to
>dismiss models of action (again, the PCT meaning), while nonPCTers tend to
>value models of action and to dismiss models of behavior.

Sorry about being so prominent, but I just can't get interested in dieting.

I don't think I can agree with the distinction you make here. Action is certainly an essential part of any control-system model -- not much would happen without it. It's not as though we have a choice between modeling action and modeling (control) behavior; to do the latter, you must do the former. In the arm model, the action is modeled as the muscle and its effect on the arm through dynamical equations.

I know that you're concerned that the fingertip in the arm model doesn't reproduce the trajectories seen in real people during rapid movements. But you have to ask whether it's important that the exact trajectories be produced -- whether it's worth the effort to adjust all the dynamic parameters so just the right movements result when a person moves the fingertip as rapidly as possible from one point to another. To focus on such uncontrolled details of behavior is to miss the main effect, which is that the model can point to a target moving at ordinary speeds in just the way a person does, and no other model yet proposed can do that. If anyone thinks it's important to match the trajectories under extreme circumstances, then quantitative experiments are called for which will show exactly how the model's movements depart from the real movements; that will show what needs to be changed in the model. I think it's premature to be worrying about such fine details; there are more important problems, such as the fact that the model starts to lose control when the fingertip gets too close to the eyes (computational problems), and can't point to a remembered target position in the dark (a major organizational problem). Solving the latter problem would make the trajectories more realistic.

My objection to putting in a lot of work to get the trajectories right is that if we were to succeed, we would gain the interest of researchers only because we have validated their concept that the output is the important factor in behavior, and we would help obscure the very point we most want to make, which is that only the EFFECT of the output really matters to the organism and is specifically controlled by the organism. The trajectories are side-effects of the details of organization; if they change, the organism would do nothing to change them back. What those trajectories accomplish are the central fact.

Even in the tracking experiments, it's possible for a person to get out a magnifying glass and point out many detailed differences between the model's handle behavior and the real handle behavior. I have had people do this. This is frustrating, because the model's movements are within five percent of the real person's movements, and often less. Talk about not seeing the forest for the trees! For the _bark_ on the trees.

Speaking of one group interesting in control:

>... the way they think that such mechanisms can be figured out is by
>>making models which show "correct" (empirically supported) actions during

>those very particular bits of control behavior. Behind this belief are,
>for example, notions in control engineering that the structure of a
>dynamic system can be identified by examining its transient (not steady-
>state) behavior.

It's true, it can, given an assumed model. But the predictions of behavior they get aren't any better than what I get in my tracking experiments with a very, very simple model -- except at the limits of task difficulty, which are seldom approached in real life. Also, there are more differences between people's behavior than there are between any test subject's behavior and the behavior of a sophisticated model matched to it. The increase in model accuracy is empty, like measuring the diameter of a soft rubber ball with a micrometer. Most important, the kinds of analysis you find in the literature again miss the forest for the bark. The concept of control of input STILL is overlooked, and as we saw last year, can be actively rejected even when it's pointed out.

>So they think that models which predict actions under precisely controlled
>circumstances are great stuff, and models which control (disregarding the
>transient details) are ho-hum.

But I say flatly that there are no such models other than the CT model, and the CT model shows that when you can predict actions accurately, you're just predicting disturbances. Show me that I'm wrong. Whose definition of "accurately" are you using? Have you got an example of an accurate non-CT prediction of actions?

>The result, I think, is that the models put forth by PCTers generally are
>not seen as very important by many non-PCT psychologists, and PCTers fail
>to see why this should be so.

I see it all right. I'm not convinced that this is a problem we're required to solve for them, or that we CAN solve for them.

Language later ...

Best,

Bill P.

Date: Sat Feb 22, 1992 1:20 pm PST
Subject: Quantitative vs. qualitative

[From bill Powers (920222.1400)]

Martin Taylor (920222.1430) --

RE: qualitative vs. quantitative thinking in psychology

I have to side with Dag Forssell on this, Martin. I understand what you mean about quantitative methods in psychology, but I don't classify these as dealing with quantitative DATA, despite the tremendous sophistication of statistical techniques. You say

>My belief is that this is so because it is working near the boundary of
>where the degrees of freedom for observation are of the same order of
>magnitude as the degrees of freedom of the system being observed. In
>physics, the systems are so simple that a human mind can more or less

>comprehend them, and relatively simple mathematics can be used to make
>excellent predictions far from the actually observed variables. In
>psychology, many different observations at different times, under
>different environmental conditions, have to be amalgamated into some
>kind of abstraction before one can see any kind of consistency.

This is putting the best face on it and what you say might well be true of your own work. However, when you consider the models used in physics, where "the systems are so simple," you find extremely complex structures in these models, extending from simple algebra through systems of hundreds of differential equations through tensor calculus. When you look at the models used in psychology, you find basically y = ax + b. Of course in order to see whether this model represents any regularity in a data set, you may have to apply very complex techniques for extracting signal from noise, but the basic model being tested is elementary, if that. So if the subject matter of psychology is so complex, why do psychologists try to handle it with such simple models?

The place where psychology is the least quantitative is in the data- taking stage. Most data exist in the form of simple and artifical events, which either occur (1) or don't occur (0). The behaviors investigated are characterized in only the crudest qualitative ways; quantitative continuous measures of behavior almost never occur except in psychophysics. Relationships are proposed only between the occurance of independent variables of a given kind and occurrance of acts of other given kinds: by and large, "dose-related" measurements are not made. Even the kinds of variables investigated, "intelligence" and so on, are defined subjectively and according to the individual investigator's private associations. The "abstractions" of which you speak are usually ill-defined, impossible to quantify, and normally handled simply by counting observations of their instances. I don't call this "quantitative."

When I read the psychology literature (a chore I mostly avoid nowadays), I see almost nothing being investigated that strikes me as a real phenomenon. Even when something real-looking is investigated, I see no quantitative measurements being made. The ONLY quantitative analysis that shows up in most articles is the statistics, which takes for granted that the data are about something and offers no explanations at all.

I think that the control-systems approach, which is fundamentally quantitative, offers the promise of handling even complex behavior in a way that is as clean as the methods of physics. I don't buy the idea that psychologists have the problems they have because of the complexity of the subject matter. I think their problems come from a primarily non-quantitative, ideosyncratic, and disorganized approach to observing human behavior, and the acceptance of very low standards for what will be considered a fact of nature. The latter bothers me the most. You can't base a science on facts that have only a 0.8 or 0.9 probability of being true. Such low-grade facts can't be put together into any kind of extended argument that requires half a dozen facts to be true at once. You need facts with probabilities of 0.9999 or better -- if you want to build an intellectual structure that will hang together. I don't think that psychology has come anywhere near meeting that requirement, individual cases aside. I would argue that we do not yet have any SCIENCE of psychology.

-- BEst,

Bill P.

Date: Sat Feb 22, 1992 3:10 pm PST
Subject: Re: Quantitative vs. qualitative

[Martin Taylor 920222 17:15]
(Bill Powers 920222.1400)

I don't want to start another side-tracking war of words that can have no benefits for CSG, but some clarification might be in order.

I argued that the subject-matter of psychology was inherently at the border of what could be experimentally investigated, and that the main problem of psychology was the focus on quantitative methods to the detriment of thought about what it was that was being measured. I said that physics was probably the only science simple enough to be comprehended now by the human mind and that great advances in physics could be achieved by the simple mathematics now available. I was misunderstood.

Bill asks:
> when you consider the models used in physics, where
>"the systems are so simple," you find extremely complex structures in these
>models, extending from simple algebra through systems of hundreds of
>differential equations through tensor calculus. When you look at the models
>used in psychology, you find basically y = ax + b. Of course in order to
>see whether this model represents any regularity in a data set, you may
>have to apply very complex techniques for extracting signal from noise, but
>the basic model being tested is elementary, if that. So if the subject
>matter of psychology is so complex, why do psychologists try to handle it
>with such simple models?

Firstly, because physics IS (barely) within the bounds of human comprehension, the brightest people tend to study physics. Do you really think that "systems of hundreds of differential equations through tensor calculus" is "extremely complex" as compared to what might be required to understand the workings of a brain (not just some simple underlying principles thereof). I would say that what you describe is simple mathematics compared to what is needed to track turbulence, even.

Why do psychologists try to handle it with such simple models? I think there is a facile answer and a subtle answer. The facile answer is that most psychologists are rather simple people, who can accept only simple models. There's some truth in it, but it does an injustice to the many who are not so simple. The more subtle answer hides in my main claim--that the complexity of psychology is near the boundary between inherently possible and inherently impossible sciences. Be careful. By "inherently" I make no reference to the intelligence of the would-be scientists. I refer to the inherent variability in what is observed relative to the possibilities of observing.

Let's look at a little analogy. A lot of the postings on the net identify perception as "what is controlled." But that is neither a correct statement of PCT (as I understand it) nor a sustainable position. Why? Because there are something on the order of 100 independently (?) adjustable joints in the human body, and many million independently excitable sensors. The sensors can react a lot faster than can any of the joint movers, perhaps 10-100 times faster in many cases. The input to the perceiving system is MILLIONS of times greater than can be controlled simultaneously. Either the input contributes to perception, or it does not. If not, what is it doing? What controls its omission from the data of perception? If yes, then it is uncontrolled perception almost entirely. It is passive. So, INHERENTLY, though behaviour can be the control of perception, perception cannot be the control of behaviour. And for exactly the same reason, not all perception can be controlled. That has nothing to do with PCT. It is in the degrees of freedom.

What can be observed in gathering scientific data is also a question of degrees of freedom, but here the difficulty is in the other sense. There are many things happening, but the possibilities for observing them are inherently limited. To observe everything would require the provision of an observing station for every degree of freedom in the universe to be observed. And some provision for accounting for disturbances from outside the universe being observed. How do we get around that? By asserting that certain symmetries exist, and doing a limited number of spot checks on the truth of that assertion. If an observation taken at time T gives result X, and another at time T+t gives result Y, what do we know? There are several possibilities, which can be grouped into three classes: The thing being observed changed, the observing device changed, or the relation between observer and observed changed. Only other observations at other times on other systems can, by analogy, lead us to bet on one of these three choices. In physics, quite a few symmetries seem to hold much of the time, and if they are carefully stated, most of the time--absolute time almost being one of them. In psychology, most of the physical symmetries are broken. The absolute time matters, both because of chemical rhythms and because of the secular aging of the subject being observed. No two people or animals are identical, and we cannot know whether a difference we observe in their actions or behaviour is due to their inherent differences or to less permanent differences. There are too many degrees of freedom for the observational tools at hand, and we haven't discovered enough about the symmetries that may exist to know whether this state will exist forever or whether it is a limitation that we may one century transcend.

What we can say is that SOMEWHERE there is an inherent boundary where the potentially observable universe exceeds the ability to observe it. There is always a practical boundary, which I place at present somewhere in the region of inorganic chemistry. That is why bright people tend to go where there exist problems that they conceivable might solve. It takes a special personality to accept the near certainty that the best they can ever achieve is a slightly closer approach to a possibly non-existent truth. So we make models that suppress almost all the dissimilarities that matter, and call the errors in our models "noise," and use the fanciest mathematics that we can justify to tease out the symmetries from the noise. And because what underlies it is so truly complex, the fancy mathematics has to be based (usually) on very simple foundations, whether it is the simple maths of control systems that may or may not be linear, whose gain can be anything great enough, whose reference weights can be quantized to plus or minus unity because the control is good enough "to fit", or whether it is the simple maths of a Clark Hull trying to fit actions to response strength gain functions, or an Estes or Mosteller fitting memorial probabilities and weights...

The more unobserved things affect what is observed, the simpler must be the descriptions of the observed, because in simplicity lies robustness, if not precision.

Enough. Here endeth the sermon.

Martin

Date: Sat Feb 22, 1992 3:37 pm PST

Subject: DEMO1 & 2 Reminder

[from Gary Cziko 920221.1945]

In Bill Powers's (920220.1800) response to the Williamses, he made mention of both his DEMO1 and DEMO2 programs.

Since a number of people have recently joined CSGnet, I want to inform them (and remind others) that both these programs are available via anonymous FTP from BIOME.BIO.NS.CA. There are both in the CSG directory and the self-extracting binary files are DEM1A.EXE and DEM2A.EXE.

These programs are also retrievable via e-mail for those of you who cannot use FTP, although this is a bit more complicated. Contact me for information concerning e-mail retrieval of thes programs.--Gary
==

INVOICE

Invoice number:_______________   (use YYMMDD)

Date:_________________________

Remit to: William T. Powers
                       73 Ridge Place, CR 510
                       Durango, CO 81301

To: _________________________   (purchaser)

____

____

____

Conditions: This invoice covers the purchase of any of three programs, named
Demo1, Demo2, and Arm, which are in your possession on approval. Payment for use of these programs by individuals for their own enlightenment is optional. For institutional or professional use, the charges (taxes included) are as follows:

Demo1: The phenomenon of control
    $35 per class-semester (or course) or other professional use
    $150 for unlimited use in teaching by a single department

Demo2: The theory of control
    $60 per class-semester (or course) or other professional use
    $150 for unlimited use in teaching by a single department.

Arm: A control-system model of pointing behavior, version 1
    $35 per class-semester (or course) or other professional use
    $150 for unlimited use in teaching by a single department.

Copies of these programs may be made freely for student use or for any other
noncommercial and nonprofessional use. The programs must be distributed as received with no changes, and may not be sold.

Program	Unit price	# courses	Total remitted
Demo1	$35	______	__________
	$150	(unlimited use)	__________
Demo2	$60	______	__________
	$150	(unlimited use)	__________
Arm	$35	______	__________
	$150	(unlimited use)	__________
		Grand total	__________

Thank you for your order.

Gary A. Cziko Telephone: (217) 333-4382

Date: Sat Feb 22, 1992 3:46 pm PST
Subject: Action Theory

[from Gary Cziko 920222.1730]

Bill Powers, Rick Marken, Wayne Hershberger, Tom Bourbon, Greg and/or Pat Williams:

I student of mine recently showed me an edited volume:

Frese, Michael., & Sabini, John. (Eds.). (1985). _Goal directed behavior: The concept of action in psychology_. Hillsdale, NJ: Erlbaum.

It contains 23 chapters, divided between Americans and Germans including two names I recognize, Gallistel and Neisser.

In the editors' introduction, they offer a definition of action theory:

"Action theory begins with a conception of human behavior: that it is directed toward the accomplishment of goals, that it is directed by plans, that those plans are hierarchically arranged, and that feedback, from the environment articulates with plans in the guidance of action." (p. xxiii)

I can see how on first glance such a theory may seem to be very similar to perceptual control theory, but I hope I am sophisticated enough by now to see where action theory (or at least this conceptualization of it) is problematic and how it is NOT the same as PCT. But I may need some help in explaining this to my student.

Any of you PCT old-timers (or anybody else) care to help out?--Gary

P.S. Does the distinction between action and behavior made recently by Greg Williams apply to the action of "action theory" as well? (I could do with some more explanation of this disctinction, Greg.)

Gary A. Cziko Telephone: (217) 333-4382

Date: Sat Feb 22, 1992 10:53 pm PST
Subject: Re: Quantitative vs. qualitative

[Rick Marken (920223)]

Martin Taylor (920222) says:

>Let's look at a little analogy. A lot of the postings on the net identify
>perception as "what is controlled." But that is neither a correct statement
>of PCT (as I understand it) nor a sustainable position. Why? Because there
>are something on the order of 100 independently (?) adjustable joints in
>the human body, and many million independently excitable sensors. The sensors
>can react a lot faster than can any of the joint movers, perhaps 10-100 times
>faster in many cases. The input to the perceiving system is MILLIONS of times
>greater than can be controlled simultaneously. Either the input contributes
>to perception, or it does not. If not, what is it doing? What controls its
>omission from the data of perception? If yes, then it is uncontrolled
>perception almost entirely. It is passive. So, INHERENTLY, though behaviour
>can be the control of perception, perception cannot be the control of
>behaviour. And for exactly the same reason, not all perception can be
>controlled. That has nothing to do with PCT. It is in the degrees of freedom

I have to admit that this is way beyond me. Are you saying that a control system does not control its perceptual signal?? Are you saying that this is the case because there might be perceptual signals that are not controlled?? I admit that I believe that a control system controls a signal which is an anolog of some variable in the system's environment. I would call this a perceptual signal. Being a control system myself I experience control as my ability to make variable aspects of what I think of as the real world (but intellectually know to be perceptual representations of an external reality) behave as I wish.
I think that this incredible realization (that behavior is the control of perception -- which can also be stated as "control systems control perceptions") is the central, incredible insight of Bill Powers (which insight is now affectionatly referred to on this net as PCT). If, as you say above, the idea that perception is "what is controlled" is neither "true nor sustainable" then I, for one, would really like to know why.
I think I can tell you why perception IS what is controlled. But that topic is adequately covered in Powers' BCP book. Please help me understand why it is not true that perception is what is controlled. Do you mean ALL perception is not controlled?? I think I can readily agree to that -- I can't control many of my perceptions (the behavior of my children being a prime example -- but the location of the Hollywood Hills is another -- just for variety). Is this what you mean? Help.

Rick

Date: Sun Feb 23, 1992 8:40 pm PST
Subject: Re: Quantitative vs. qualitative

[Martin Taylor 920223 15:15]
(Rick Marken 920223)

Oh. dear. How easily we are misunderstood!

I said:
>>INHERENTLY, though behaviour
>>can be the control of perception, perception cannot be the control of
>>behaviour. And for exactly the same reason, not all perception can be
>>controlled. That has nothing to do with PCT. It is in the degrees of freedom

and Rick interpreted me as saying

>Are you saying that a control system does not control its perceptual signal??

No, exactly the opposite. I am saying that Bill's thesis can be true (and on
other grounds must be true), but that the S-R position CANNOT be true.

What I was trying to point out was that not all of perception can be
simultaneously the object of control. It was only a little reminder to
the many posters to the net who carelessly write as if the fact that
all behaviour is the control of some perception is the same as all perception
is under the control of some behaviour. Bill never writes that way, and
even in this posting you support me.

It is perhaps worth following up on this. I meant to do it in another way,
later, because I am very busy right now, but here goes.

J.G.Taylor claimed that all perception was the result of the ability to control
it (not the fact of controlling it, which was a misapprehension Rick had a
few months ago). I wrote a critique of this position in 1973, on much the
same grounds as the one to which Rick commented today. The solution in 1973
was in part the same as the solution today, which is to say that much of the
perceptual structure is passive rather than active. It is based on the high
statistical correlation of structures and events to which the sensors are
exposed. The information coming into the sensor array is highly redundant,
which means that it can be reduced by fairly simple tricks into a more nearly
orthogonal space. I don't know how much reduction can be safely accomplished
passively, but I should be surprised if it isn't 3 orders of magnitude.

If JG was right, then what happens is the building of many control systems,
each capable of handling some component of the remaining incoming information,
but not themselves orthogonal in what they control. There will still be some
3 orders of magnitude difference between the number of potentially controllable
perceptual abstractions at any moment and the number of degrees of freedom
avaialble to control them. So, either only a few control systems are actually
controlling at any moment, or the many that are controlling are doing so in
a distributed, coordinated way. This has nothing to do with any hierarchy
of control, but applies equally if you want to consider the hierarchy.

A control system that is not controlling is not the same as one whose reference
signal is zero. But it is the same as one whose gain is zero. So, whether
there are only a few systems actually exerting control at any moment, out of
the many that could do so, or whether a large proportion are controlling
in a coordinated way, there will be many in which the percept does not match
the reference even when the whole system has come to a stable state. In fact,
almost all of the ECSs will fail to match percept to reference. This is not
a failure of the system, but is inherent in the notion that there is more that
could be controlled for than can be handled at any moment.

Imagine a situation in which stability and correct control has been achieved in some part of the hierarchy. By definition, the error signals in the stable part are zero. This means that the reference signals fed into all but possibly the highest ECS in the stable part are also zero, and in turn this means that the percepts in that part are also zero. In other words, the perceptions in a stabilized hierarchy are (with the exception of possibly predetermined and non-varying biases) all empty. To take a mundane example, one does not feel the gusts of crosswind one expects and controls the car against.

Conscious perception (and presumably unconscious perception as well, since we have no real idea what either means) is of those things that are not being currently controlled for exactly. But it is of those things for which we might control, if JG is correct.

I think that this view is, in a simple-minded way, not far from being right. But the probable existence of coordinated distributed control means that even when perceptions are being controlled optimally, there are many reference signals that are not matched by their percepts. That mismatch represents a tension among the possible (but not presently actual) states of the world. To refer back to another long-ago discussion, it is one way of looking at subjective probability from the viewpoint of an outside observer (illegitimate, perhaps, but there is a connection between the tolerance the ECS will have for mismatch and the expected degree of mismatch when perfect distributed control is achieved).

In summary, if the perceptual degrees of freedom are higher than the action degrees of freedom (and they are MUCH higher in practice), then perception cannot control behaviour, but all behaviour can be the control of some perception. If the large number of perceptual degrees of freedom have any evolutionary relevance, they imply that there are things in the environment that might affect them. Those things could be evolutionarily relevant only if there were somne conceivable actions that could affect the associated percepts, and hence affecting those things could be part of a control loop controlling some perception. That in turn would mean that there are many more ECSs than can be simultaneously satisfied, so that there are necessarily non-zero percepts at all times, regardless of the state of stability of the control hierarchy.

I intended a demonstration that behaviour is the control of perception, rather than the reverse, and go swatted for it by the loose canon of the church. I trust I have made amends.

Galileo

... And yet non-zero percepts exist.

Date: Sun Feb 23, 1992 8:40 pm PST
Subject: LFG

Re: Bruce Nevin (Fri 920121 07:36:42)]

>This necessarily reduces PSG-based syntax to a filtering function
>ratifying an unstructured collocation of words motivated by nonverbal
>perceptions. Just throw out words in whatever order they occur to yo

>...

This isn't true, basically because processing activities can be interleaved in complicated ways: one of the guiding principles of LFG was that the results of proccessing in accordance with the grammar should be highly independent of the order in which processing operations are carried out. Dependency-grammar-based generation might be smoother than LFG-based generation, but I don't think it's worth arguing about this in the absence of actual implementations that people can play with.

>Whatever subclasses of words you define--human, animate,
>abstract, concrete, etc.--as soon as you say only +animate nouns can
>cooccur with a +animate verb, someone will think of a perceptual
>situation, in describing which you would use a -animate noun with a
>+animate verb. The teacup danced off the table.

I'm not at all sure what you're getting at here. Since the late sixties, generative syntacticians have mostly assumed that selection restrictions were not directly syntactic, but due to syntax-semantics interactions. If my LFG system had a real semantics associated with it, it would detect that `the teacup danced off the table' is peculiar on the basis of attributing contradictory properties to the teacup (it doesn't now, because its `semantics' is just a display facility, without any kind of inferencing system). But I don't see what relevance this has to how the syntax itself is organized.

>To address the same word-word : perception-perception association in a
>PSG-based system like LFG requires an inordinately complex
>superstructure of Principles and Constraints and other Generalizations.
>Much of the hoopla in the literature takes this form:

Well, my LFG system fits within the 640K boundaries of an XT, which makes it in a sense simpler than any other that I've heard of. In fact the *only* other PC-based parsing system for an actual linguistic theory that I know of is also for LFG, done at the University of Stuttgart in 1986. This is a prima facie argument that LFG is pretty simple.

Keep in mind that the apparatus of PS rules & grammatical functions is there to make it possible to deal with the existence of a variety of ways of expressing the arguments of predicates. Some language use linear order, others, case-marking, yet others agreement, plus assorted combinations. Some of the possibilities are presented in the grammar fragments that come with the system.

In spite of the above, I do agree that there is rather too much emphasis in current generative grammar on `Principle Hacking' of the kind that Bruce derides, and not enough on finding ways to produce clear & well-defined analyses of novel grammatical phenomena.  I got attracted to LFG because it seemed to do better than other frameworks on some of the topics I was working on at the time (`long distance' agreement, case-marking, & morphological blocking; highlights of analyses found in the Icelandic, Irish & Spanish fragments). There is *lots* of room for improvement, but I'm not actually convinced that dependency grammar would do significantly better on the stuff.

My belief is that things would be in a much healthier state if people paid more attention to how well (or more realistically, how badly) their analyses actually worked, and I suppose I could present my package as an effort to shift things a bit in that direction. It is in particular strictly for *introductory* students. As things stand, people get introduced to syntax via TG, in which it is virtually impossible to write analyses that work. The ultimate consequences of this is that people get habituated to all sorts of vagueness & omission, in a way which was maybe unavoidable 20 years ago, but is nowadays quite indefensible.

Avery.Andrews@anu.edu.au

Date: Sun Feb 23, 1992 8:40 pm PST
Subject: From Dick Robertson before going to Belgium & on reorganiz.

[From Dick Robertson]
There has been so much exciting stuff on the net lately, I'm sorry I couldn't put in my 2 cents worth, but I've been preparing 10 lectures on PCT to present to a group of clinical psychology students at the univ of Louvain la Neuve. Anybody got any last minute suggestions?

(to Rick Marken)

Thanks for the kind words on the old Powers Game research. If any body wants to look at what we reported there on reorganization in learning the reference is Perceptual and Motor Skills, 1985, 61, 55-64, The Phantom Plateau Returns.

(to Francis Heylighen)

I saw your note about your funding problem, and will give you a call from Louvain la Neuve, and write that letter , if it might help, although your government sure never heard of me.

(to Bill Powers)

I guess you never got that post I tried to send you about the grade control stuff. I'll try again when I'm back in May.

Best Wishes to everybody.

Dick Robertson urrobert@UXA.ECN.BGU.EDU

Date: Mon Feb 24, 1992 2:53 pm PST
Subject: Re: bridges

From Tom Bourbon [920224 -- 12:56] [If my connection holds up for a few minutes -- this is my first connection with the net in 4 days.] Greg Wiliams [920221 -- 17:11]. Aha! to your aha! The few comments I was able to post during the pre-cease-fire debate on Beer did not work as I intended -- I did not draw the distinction you said came

through in your epiphay. The emphasis on perception, not action, as the topic of interest is the greatest difference between PCTers and non-PCTers. What some people took as dogma on the part of PCTers, during the debate, was often a declaration that Beer's work contains no clear indication that he is interested in whether, or how, organisms control perceptions. As Rick said several times, if Beer does do that, I am willing to be convinced.
****
5% solution. While I was off the net, Rick and Bill accurately summarized some of my data which figure into the (actually better than) "prediction to within 5%" remarks by Bill.
****
On modeling and robots. There may already be a post about this, but there is no guarantee I will get to it today. In the latest issue of *Psychological Science* (vol. 3, no. 1) is an article that might be of interest to people on CSG-L:

Teitelbaum, P. & Pellis, S. M. (1992). Toward a synthetic physiological psychology. pp. 4-20.

The authors have some interesting things to say concerning the relationships among reduction and reductionism, which they construe as not identical, and about physiological details and functional analyses of behavior. I believe their comments are germane to the recent debate about physiological realism on csg-l.

Even more interesting to me is their call for testing ideas about behavior, not through the hypothetico-deductive method and its accompanying statistical procedures, but through the construction of robots. To quote:

"Therefore, to build psychology rather than medicine, synthetic physiological psychology aims to create a physically achieved analytic approach to the design not of people but of robots that behave like people. That is the first meaning of the word synthetic: artificial, man-made, as in a robot. Our concern, of course, is with software, not hardware." p. 6.

I offer this, not as a spark to rekindle the debate, but as a nice convergence of approaches upon the method of modeling. And, of course, there is the obvious, direct tie to Bill Powers' articles on the design of robots, in *Byte*, in 1974. Teitelbaum and Pellis do not cite *Byte*, (a letter to the editor is on the way, mentioning that precedent), but they do briefly mention Powers' *Behavior: The control of perception*. Good reading.

Tom Bourbon <TBourbon@SFAustin.BitNet>
Dept. of Psychology
Stephen F. Austin State Univ.
Nacogdoches, TX 75962 Ph. (409)568-4402

Date: Mon Feb 24, 1992 2:56 pm PST
Subject: Linguistics; scientific momentum

[From Bill Powers (920224.0800)]

Bruce Nevin (920221) & Avery Andrews (920223 & prev) --

RE: LFG, dependency grammar, and momentum.

What's missing for me from all these approaches to grammar is the sense of

"Yeah, that's how I do it." It could be claimed that the mental machinery is invisible to the observer and that all we get out of it is the result. But it seems strange to me that these methods use words, symbols, and rules for manipulating them in the familiar way, except for the fact that they propose content that is strange to me -- Bruce's expansions, and Avery's NP, VP, trees, recursions, etc.

The basic concept of control theory, particularly control of perception, does strike people who internalize the concept as "Yeah, that's the way I work." There is still a lot of hidden machinery, and the closed-loop relationships aren't always intuitively obvious, but enough of the iceberg seems to show to be recognizeable to the occupant of said brain. In the case of linguistic models, the part of the iceberg that shows in the description doesn't seem to match any experience I can find in here.

I've tried for quite some time now (longer with Bruce) to elicit a description of what a linguist is doing in the process of getting from a received sentence to the structural analysis. What I get back is further analysis -- i.e., you DO it, but you don't DESCRIBE WHAT YOU'RE DOING. I don't want to know that "bite" is a word that takes two arguments, or that "bite" is a PRED function of SUBJ,OBJ. That doesn't tell me what you're doing in your head to get from sentences to those statements. I don't want to know what a system or theory says about how the words are related -- what I'm trying to get is a description of what happens in your consciousness when you begin with a new sentence (Josh felt Jean was unsympathetic) and begin to re-represent it or something about it. Instead of arguments developed from further application of the method in question, I'm trying to hear the processes going on in present time in the person offering those arguments.

In short, I'm asking you to get out of the "i am a linguist" mode and go up a level. Get outside of the game. Break the set. Reorganize the Gestalt. SLOW DOWN. Look at what you're doing. Stop being an expert. Whatever.

Some time ago I offered an amateur proposition about how we construct sentences. It had to do with describing meanings. You pick a meaning to be mentioned first, mention it, and then tack on any refinements, resolve ambiguities, and add details that seem necessary as you contemplate the words so far, and when you're satisfied go on to mention the next meaning, and so on until you feel that the sentence, when heard, will more or less point to what you meant to communicate. I could feel myself doing that as I wrote the previous sentence, not knowing how it would end when I started it. I'll bet you can feel yourself doing it when you construct sentences, too. I will also bet that while you're doing that you aren't thinking "Gee, this word requires two arguments, one of which I've already mentioned, and the other of which could be any one of 12,000 words in my lexicon." And that you aren't building up a huge LFG diagram, either. I'll bet you do it just the way I do, with perhaps some specialized linguist-type commentary in the background.

My proposition is full of unanswered questions. How do we "mention a meaning?" How do we "tack on refinements" and "resolve ambiguities?" How do we know which previous subject is referred to by a later word? What conventions are there that help us do all this in a way that someone else will be able to unpack? For all I know, some current linguistic theory will have something useful to say on those subjects. But whatever it has to say, I want to get the feeling, upon hearing it, that "Yeah, that's how it works." I want to be able to look more closely, or less closely, at myself making sentences and RECOGNIZE what is being proposed. RECOGNIZING the process to which the theoretical description refers is what gives us the feeling of verification. We NOTICE IT HAPPENING, not hypothetically but actually, in present-time experience. It's like noticing that whatever you're controlling, it's a perception. This proposition isn't verified by giving logical arguments in favor of it, according to someone's theory. It's verified by seeing it going on.

I'm suggesting strongly that answers to questions like those at the start of the previous paragraph are not to be answered by reference to a theory, but by reference to direct experience. Only when you run out of information of that kind is it time to start proposing hidden machinery. I feel that there is a hell of a lot going on while we make language that we CAN notice; we are nowhere near exhausting the realm of what is available to consciousness -- provided that consciousness doesn't get stuck at one level of observation. One way to get stuck is to become committed to a logical system in preference to further investigation of direct experience.

Which brings me to the third RE: momentum.

I'm more or less resigned to the fact that when people from other disciplines get interested in control theory, they already have built up a lot of scientific momentum. They are not starting from scratch. They have built up a complex structure of understandings in the course of trying to make sense of whatever aspects of life seem most interesting. They generally think they've been getting somewhere, although an interest in control theory shows that they see some unsolved problems. What takes a long time to realize -- years -- is that if control theory is a correct description of how people work, then way back in the mists of the past, all the ideas they have built upon for all these years contain some fatally wrong assumptions. Way back down there in the foundations.

It's very difficult to trace an adult and well-developed world view back to its origins. Usually this comes about when some belief begins to look shaky and you try to shore it up. You say, "Well, this is true because of A, B, and C," and if there's still a shred of puzzlement you try to backtrack on A, then B, then C: how did I get to them? But if you have control theory in your head, it's going to clash, sooner or later, with one of those antecedent ideas, because they're sure to contain some cause-effect notion that's incompatible with negative feedback and its implications. Bingo, a conflict.

If the conflict is resolved in favor of CT, then the antecent idea is marked "wrong." But if it's wrong, then something is wrong with the next layer that depends on it, and if something's wrong at that layer, the next layer is off the track, too, and now you have a BIG conflict, because the belief that looked slightly shaky when you started this has lost its foundations and is threatening to collapse.

So what do you do then? I think that the first thing you do is look for a way to make it right again. To hell with the antecent propositions, it's GOT to be right. I've committed too much of my life to it to just junk it. It would be too embarrassing (thinking of all the thousands of students hanging on your words and writing them down, all your publications, all your honors, all the arguments you've won). This kind of conflict can wreck your professional life.

If you don't want that kind of conflict, stay away from revolutionary ideas, particularly those that sucker you into thinking they're probably right. Control theory has shattered more than one career. We've had a few graduate students on this net whom I advised to wait until they had their degrees to take up control theory. I'm sure the ones who didn't take my advice have been paying the price: how can you satisfy your committee when you no longer believe anything they would like to hear you saying?

The questions I'm asking of the linguists are similar to those I'd like to ask everyone on this net, in their own fields. It simply stands to reason: control theory was not around when your discipline got its start, and it wasn't around while that discipline grew from an idea in somebody's head into an elaborate series of deductions from basic premises and observations interpreted in the light of those premises. If those premises didn't include the principles of control, they were probably wrong. If they were wrong, they are still wrong, and it's unlikely that very much of the rest that depends on them is right. If you don't want to track down the wrongnesses and revise your discipline from its foundations, then there's really no point in your going on with control theory. If you just dabble in it or use it to show how right you've been all along, nothing much will come of it, for you. Control theory is a revolutionary idea. What you do with revolutionary ideas is have a revolution.

Best to all,

Bill P.

Date: Mon Feb 24, 1992 3:52 pm PST
Subject: Old business

From Pat & Greg Williams (920224)

In a post dated 920220.0900, Bill Powers wrote the following:

P>I believe that if you read what I wrote more carefully, you'll see that I
P>didn't claim that every control-system model I've proposed matches human
P>behavior within 5 per cent. This was intended to apply only to those simple
P>models we have been able to test quantitatively against human behavioral
P>data. If I didn't say that, I certainly should have. I usually remember to
P>include that limitation.

That was in reply to this question of ours in a post dated 920219:

W>Question: Where does the 5% (or better) accuracy figure for PCT models come
W>from? Which models, what data? I (Greg) admit to surprise regarding this
W>figure for the Little Man, as I didn't know Bill had compared it to "real"
W>data.

Here are the two posts (the first from Bill, dated 920218.0900, and the second, from Rick Marken, dated 920219) which resulted in our raising the question:

P>Much of the verifiability of my model is behavioral. In tasks ranging from
P>elementary to somewhat complex, my model generates predicted traces of limb
P>movements that match actual limb movements with an error of five percent or
P>better.

M>PCT models behave, too. In fact, they control. And they mimic human control
M>(as Bill mentioned) to within 5 per cent (far better in most cases). These
M>are not statistical matches -- the model produces behavior that matches
M>human behavior perfectly every time.

In the same post (920220.0900) in which Bill said that the "5 per cent [figure] ... was intended to apply only to those simple models...", Bill says:

P>This [5%] figure refers to models of visual-motor tracking performance,
P>where we have accurate comparisons of model and human data. I have done these
P>comparisons using many unpredictable disturbances, changes in handle
P>sensitivity, nonlinear connections from handle to cursor, and multiple
P>disturbances acting at different places on the display.

and

P>The Little Man hasn't been compared with quantitative human data (by me).
P>It does predict some aspects of pointing behavior quite accurately: the
P>finger follows a slowly-moving target just as a real person's would -- the
P>tracking error in both cases is essentially zero, and is certainly less
P>than 5 per cent of the range of target movement in three dimensions. An
P>open-loop model of pointing behavior published in Science a few years ago
P>achieved a minimum pointing error of something like 4 degrees in direction
P>and 4 centimeters in position, if I remember correctly -- with a STATIONARY
P>target.

Bill, we DO try to read your posts carefully. But I think you can see from the above that you need to read ours carefully, too. We did not imply that you had claimed that every control system you had proposed matches human behavior within 5%. Greg did express wonder at the 5% accuracy for the Little Man, simply because of your 920218.0900 post. In your 920220.0900 post, you retract that. Next time you're tempted to accuse us of not reading your posts carefully, we advise you to read ours AND your own more carefully.

Re. LCS II: We still need answers to questions posted last week. Your disk was D.O.A. -- was a Preface on it???? We'll type in the two papers here, but please send the Preface again (?). Tom's Foreword galley goes to him tomorrow.

Pat & Greg

Date: Mon Feb 24, 1992 4:07 pm PST
Subject: linguistics

Re: Bill Powers (920224.0800)

>What's missing for me from all these approaches to grammar is the sense of
>"Yeah, that's how I do it." It could be claimed that the mental machinery
>...

True, but I think this would be true of *any* theory of higher order perception, which is what Chomskyan linguistics should be seen as. After all (if I read and recall my secondary sources correctly), it was in trying to introspect on how `sensations' were `apperceived' into `perceptions' [?], e.g. how points of color in the visual field were integrated into the perception of, say, an apple, that pre-behaviorist psychology went way wrong.

Linguists typically believe that people can introspect about what sentences are grammatical & what they mean (similar to what perceputal psychologists do with necker cubes, etc.), but that they have no insight whatever about how they know these things.

>I've tried for quite some time now (longer with Bruce) to elicit a
>description of what a linguist is doing in the process of getting from a
>received sentence to the structural analysis.

My belief is that there's just a dumb unit (sort of like a Fodorian `module') that churns away trying to assign grammatical structures to input, similar in general nature, and perhaps actually based on, whatever it is that resolves the visual field into objects of various categories, bearing assorted properties and relations to each other. Both grammatical structure and the consequent meanings, and `visual structure' are assigned (to streams of noise and the visual field, respectively) subject to the constraint that the whole scene has to make sense, so that, for example, we sometimes have trouble recognizing people in places we don't expect to see them.

So my grammars, parsers, etc., are just tentative attempts to get a grip on what this unit is doing. Given the failure of pre-behaviorist psychology, I'm not at all inclined to abandon this aspect of standard linguistics.

As for `momentum', it strikes me that one place where PCT has not made much progress is in higher-level perception. E.g., noone knows how to build a `fridge-door open' recognizer. The work I know of that has actually gone anywhere in this direction is `cognitivist', e.g. Marr, Pinker, Jackendoff, etc. But standard cognitivism doesn't present a satisfying (to my eye) picture of what the higher order perceptions (`mental representations'), are actually doing - when it comes to saying what makes a representation `true', all anyone has to offer is idiotic `correspondence theories', followed by silence. Here, I think, PCT has a definite & important contribution to make, since one can say that the H.O.P.'s are there to be controlled (and thereby satisfy the intrinsic reference levels, and incidentally keep the organism alive).

So my inclination is to take the current theories of H.O.P. more or less as is, & see what can be gotten by adding control to the mixture. & I guess I ought to add that it wouldn't bother me to find a few S-R sub-systems lurking around in the corners. One good candidate is the one that directs your attention to any bit of speech you hear in which something that sounds like your name is mentioned. Another would be the one that makes it so wierd to walk up or down a stalled escalator (I guess what's going on here is that there's a circuit that correlates locomotion rate with flow of the visual field, and when something tells it that one is on an escalator, the expected correlation is altered).

Avery.Andrews@anu.edu.au

Date: Mon Feb 24, 1992 6:42 pm PST
Subject: bridges again

From Greg Williams (920224)

>Rick Marken (920221)

>Halleluja. [etc.]

I was not quite prepared for the resonance (a Sixties' word?) which Rick -- a battle-hardened PCT modeler who has been on the front lines skirmishing with (definitely!) non-PCT psychologists -- expressed regarding my hypothesis about problems due to the contrasting importance of (PCT-definition) action and (PCT-definition) behavior as viewed by PCTers and nonPCTers. Perhaps I am on to something. But Rick is on to more; he is way ahead of me with a proposed remedy, preparing to meet the opposition head-on, as it were, with demonstrations of how simple behaviors can be accomplished by various and complicated sorts of actions. I myself still think Bill Powers' "Feedback Model for Behavior" article (1971) is quite convincing that at least some kinds of actions reflect disturbances virtually 100% -- in other words, in these cases, the organism is achieving excellent control, with little (excuse the term) slop. I think I agree with Bill that this is the BIG thing which PCTers have to teach nonPCT psychologists: there indeed are cases when measuring actions will tell much about non-organismic factors (in a word, disturbances) and little about the control structures of the excellently controlling organisms. Non-PCT psychologists need to understand that they will not always learn much about an organism by precisely measuring its actions, other than that it is controlling quite well, thank you. Rick's proposed project certainly might help get this notion across.

In addition to the approach of trying to moderate nonPCT psychologists' excessive enthusiasm for the explanatory ability of actions, there is a second approach, also being followed by PCTers to an extent, as has been noted recently by Bill Powers. This approach attempts to show nonPCT psychologists that PCT-based models can do what (some) nonPCT psychologists want to see: match data on organismic actions. By "match," I don't necessarily mean exact quantitative prediction of data by the models. But I think what will really "wow" some of the skeptics will be combinations of models and data which allow choices to be made between various fairly detailed control mechanisms. I don't think this can be done to any significant extent unless transient ("poor" control) situations are investigated; steady-state ("good" control) dynamics almost entirely mirror changes in disturbances, as noted above. Here, "poor" control could be "difficult" control, but it also could be a "start-up" situation with "easy" control. The important requirement is that control not be anywhere near perfect, since then the data will tell little about the "internal" dynamics, and one cannot easily distinguish between control structures or determine the parameters of such structures.

Even if many nonPCT psychologists were convinced (by the first approach discussed above) of the truth that actions often (when control is good) tell us little about organisms and a lot about their environments, at least some want to learn about the mechanisms of control -- and they are ripe for the second approach.

In short, I think both approaches are important and, to the extent possible, given the limited time and resources of PCTers, should be given high priorities.

Greg

Date: Tue Feb 25, 1992 5:39 am PST
Subject: CACS92

>From Agnes Guillot (920225)

Call for Participation

COMPARATIVE APPROACHES TO COGNITIVE SCIENCE (CACS92):

An International Summer School

CACS92 is an international summer school to be held in Aix en Provence, France, July 6-17, 1992 on comparative approaches to cognitive science. This school will bring together leading investigators in animal and human cognition, artificial intelligence, and robotics to discuss, compare, and share the concepts, problems, and techniques that characterize their fields of investigation. It will also offer numerous opportunities for collaboration. Its main goals are to discuss the role that investigations of animals and machines can play in the development of cognitive science generally, and to provide the intellectual and methodological tools necessary to the advancement of such developments.

A major focus in cognitive science has been on modeling the performance of tasks that are characteristic of human intelligence, such as planning, problem solving, scientific creativity, and the like. Several investigators have recently suggested the possibility of a complementary comparative approach to cognitive science. Rather than modeling toy problems from the larger domain of human expert behavior, this approach advocates the modeling of whole, albeit simple, organisms in a real environment, performing real biological tasks (surviving, exploring, mating, feeding, escaping predators, etc.). The goal of the approach is to develop coherent incremental models out of functionally complete components. Achieving this goal requires that we investigate animal performance and the mechanisms they use as the basis for our growing models. It also requires extensive collaborations among ethologists, psychologists, computer scientists, engineers, and cognitive scientists because no one of these fields, by itself, has the tools to thoroughly understand the mechanisms of such complex processes. The purpose of this summer school is to review the state of the art in this interdisciplinary approach and to share the tools and perspectives it requires.

The summer school will be held at the Ecole d'Art d'Aix en Provence. Aix is in a beautiful part of France known as a favorite location for many of the Impressionist painters.

The summer school will consist of morning lectures followed by afternoon discussions. English will be the official language. We have asked the instructors to prepare presentations that are accessible as tutorials to the students and are broader than normal, that describe not only the investigator's own interests, but also review the state of the art, and describe the theoretical and empirical tools that are employed. We have also asked them to draw explicit conclusions concerning how the work they describe impacts on cognitive science more generally. Presentations will draw specific conclusions about the role that cognition plays in solving behavioral problems and identify the kinds of organisms and environments in which such mechanisms may be useful.

During the summer school, the Ecole d'Art d'Aix will simultaneously organize a series of artistic activities and demonstrations including conferences, workshops, and shows, for which artists of many nationalities have been invited to contribute works along themes related to those of the planned summer school (artificial life, behavioral organization, networks, interconnectedness, robots, animal behavior, etc.). Many opportunities for interaction among the scientific and artistic participants will be available.

This promises to be an excellent and influential summer school. In addition to the invited speakers, a limited number of participants/students can be accommodated. Advanced graduate students, young researchers, new PhDs, and post-docs are particularly welcome. Participants are invited to submit abstracts for poster presentations during the summer school.

The costs to participants have not yet been determined. We expect that the registration fee for the summer school will be approximately FF4,500, which would cover summer school registration, room (in student housing at the University of Aix) and board. We expect some scholarship support to be available to help offset these costs.

Prospective participants are urged to indicate their interest as soon as possible because space is limited. Participants should submit the following: A letter describing their interest in the subject matter of the conference and a curriculum vitae. Include a full mailing address, electronic mail address, and FAX number. If scholarship support is desired then a letter of recommendation from the participant's advisor or department chair is also required. Please indicate the amount of scholarship support desired. Those desiring to present posters should submit a one-page abstract. Centered at the top of the page should be the complete title, author name(s) with the presenting author underlined, affiliation(s), and complete mailing address. This is followed by a blank space and the text of the abstract.

One copy of all material should be sent to each of the summer school organizers:

Herbert ROITBLAT
Department of Psychology
University of Hawaii at Manoa
2430 Campus Road
Honolulu, HI 96822
USA
email: roitblat@uhunix.bitnet
roitblat@uhunix.uhcc.hawaii.edu

Jean-Arcady MEYER
Groupe de Bioinformatique
URA686. Ecole Normale Superieure
46 rue d'Ulm
75230 Paris Cedex 05
France
e-mail: meyer@wotan.ens.fr
meyer@frulm63.bitnet

Tentative Program

Comparative Approaches to Cognitive Science (CACS92)

INTRODUCTION

Jean-Pierre Changeux (France) From non-human to human cognition: challenge and prospects

Herbert Roitblat (USA) Comparative approach as a tool in cognitive science

Jean-Arcady Meyer (France) Computational approaches to cognition

Marc Bekoff (USA) Cognitive ethology, common sense, and the explanation of animal behavior

PERCEPTION AND ACTION

Tom Bourbon (USA) Perceptual control theory: Modelling conflict, cooperation and control

George Butterworth (UK) Factors in visual attention eliciting manual pointing in human infancy

Steven Whitehead (USA) Towards a computational theory of perception, action and learning

CONCEPT FORMATION

Roger Thompson (USA) Natural concepts and self-concept in animals

Lorenzo Von Fersen (Germany) Abstract and natural concept formation in animals

Keith Holyoak (USA) Natural and artificial induction

INTERNAL WORLD MODELS

Julie Neiworth (USA) Internal models of space, time, and movement in animals

Catherine Thinus-Blanc (France) Spatial information processing in animals

Bartlett Mel (USA) Mechanisms and applications of associative learning in biological sensory and motor systems

MOTIVATION AND EMOTION

Frederick Toates (UK) Animal motivation and cognition

Janet Halperin (Canada) Cognition and emotion in animals and machines

Niko Frijda (Netherlands) Emotions in robots

INTENTIONALITY

Daniel Dennett (USA) Animals and human beings as intentional systems: The fundamental difference

David McFarland(UK) Goals, no-goals and own-goals

Peter Kugler (USA) Informational fields and intentional action

Colin Allen (USA) Intentionality: natural and artificial

LANGUAGE, COMMUNICATION AND COOPERATIVE BEHAVIOR

Peter Marler (USA) Communication in animals

Sue Savage-Rumbaugh (USA) Cooperative communication by pygmy chimpanzees

Giulio Sandini (Italy) Cellular robotic systems

LEARNING

Randy Gallistel (USA) Time representation and conditioning in animals

Jean Delacour (France) The memory system of the mammalian brain

Richard Sutton (USA) Learning and planning

Leslie Kaelbling (USA) Reinforcement learning in robots

EVOLUTION

David Premack (France) Evolution of cognition from primates to man

Michael Dyer (USA) Symbol Grounding and Evolution of Primitive Communication

Rik Belew (USA) Interacting models of evolution, ontogeny and learning

CONCLUSIONS

John Anderson (USA) Relationship between rational analysis and
cognitive mechanism

Margaret Boden (UK) Creativity in Humans and Machines

Rodney Brooks (USA) Prospects for cognition in autonomous robots

END

Date: Tue Feb 25, 1992 9:54 am PST
Subject: Re: Quantitative vs. qualitative

[From Rick Marken (920223)]
Martin Taylor (920223) says:

>What I was trying to point out was that not all of perception can be
>simultaneously the object of control. It was only a little reminder to
>the many posters to the net who carelessly write as if the fact that
>all behaviour is the control of some perception is the same as all perception
>is under the control of some behaviour. Bill never writes that way, and
>even in this posting you support me.

Worthwile point -- I agree.

> There will still be some
>3 orders of magnitude difference between the number of potentially controllable
>perceptual abstractions at any moment and the number of degrees of freedom
>avaialble to control them.

How could the number of "potentially controllable" perceptions matter?
How could such a count even be determined?

>So, either only a few control systems are actually
>controlling at any moment, or the many that are controlling are doing so in
>a distributed, coordinated way. This has nothing to do with any hierarchy
>of control, but applies equally if you want to consider the hierarchy.

I don't understand this at all.

>A control system that is not controlling is not the same as one whose reference
>signal is zero.

How?

>But it is the same as one whose gain is zero.

Why can't it be either one?

>So, whether there are only a few systems actually exerting control at any >moment, out of the many that could do so, or whether a large proportion are >controlling in a coordinated way, there will be many in which the percept does >not match the reference even when the whole system has come to a stable state.
>In fact, almost all of the ECSs will fail to match percept to reference.

This is not true, as you can see in my hierarchical model. If set up properly (so that the perceptions are relatively orthogonal and there are sufficient orthogonal outputs to control them, all systems bring perceptions to their references; the error may not always be precisely zero but it remains within a very small percent of its possible range.

>This is not
>a failure of the system, but is inherent in the notion that there is more that
>could be controlled for than can be handled at any moment.

This is again not clear. What is this "could be controlled" stuff? Is it existing control systems that could be used but, if they were, would lead to possible conflict due lack of sufficient independent output degrees of freedom? If so, OK. But if these are just perceptual variables "out there" that "could be controlled" then I think they are irrelevant; they just don't exist from the point of view of the system -- because there are no control systems around to control them.

>Imagine a situation in which stability and correct control has been achieved
>in some part of the hierarchy. By definition, the error signals in the
>stable part are zero.

Not true.

>This means that the reference signals fed into all
>but possibly the highest ECS in the stable part are also zero,

Not true. Reference signals can be the result of integrations, for example. Look at the spreadsheet!

>and in turn
>this means that the percepts in that part are also zero. In other words,
>the perceptions in a stabilized hierarchy are (with the exception of
>possibly predetermined and non-varying biases) all empty.

That would be a strange experience!!

>To take a mundane example, one does not feel the gusts of crosswind one expects >and controls the car against.

One NEVER does -- just their effects on a controlled variable -- position of the car -- which can be kept at a non-zero reference.

> But it is of those things for which we
>might control, if JG is correct.

I'm still not sure what JG is saying, but if it is any of the above then I'm afraid JG is wrong.

To Tom Bourbon:

I see you are listed as one of the presenters at that conference in France. I want to go too!!!! Aix-en-provence -- give me a break. How come they have all those know-nothings talking about intentionality and I'm not on the list??? I am truly envious and proud of you -- give 'em hell.

Regards

Rick

Date: Tue Feb 25, 1992 10:07 am PST
Subject: 5%; BEERBUG

[From Bill Powers (920225.0800)]

RE: legal matters:

G & P Williams (920224) --

When I referred to
P> ... predicted traces of limb
P> movements that match actual limb movements with an error of five percent
P> or better.

I was thinking of tracking experiments (in which handle movements are a function of limb movements, which can be directly deduced from behavioral recordings if you know shoulder position relative to handle position and length of handle stick and arm segments). As you know, I don't have any way to measure behavior in a way pertinent to the Little Man model, and I don't lie about things like that to my friends. Even if I did lie, it wasn't done on my computer. And if it was done on my computer, I wasn't home at the time. But my lawyer says not to admit anything.

I came across an interesting fact about cockroaches, in Beer's book, p. 82, Fig. 4.5. Here a real cockroach's gaits from slow to fast are shown as a plot of leg down- and up-time for each leg. The patterns look very confusing at first, until you realize that both left and right sides follow exactly the same forward-swinging pattern regardless of speed of walking, except for the rate at which the left and right patterns alternate.

In a plot for a real bug, the forward swings (foot-up pattern) for the left side look like this at the slowest speed (the spacings are as near to quantitative as I can reproduce by eye and ASCII):

```
                  _           _           _           (rear)
                    _           _           _
                      _           _           _       (front)
```

And for the right side they look like this:

```
                        _           _           _
                          _           _           _
                            _           _           _
```

This is not a plot of foot-placements but of durations of up-time (dashes) and down-time (spaces). At a high speed the two patterns look like this:

```
Left:             _   _   _                 (rear)
                    _   _   _
                      _   _   _             (front)
Right:              _   _   _
                      _   _   _
                        _   _   _
```

Note: The rear-middle-front sequence on each side takes exactly the same length of time at either speed, and the duration of foot-up time is the same at any speed. So all that changes is the duration of the foot-down period (spaces) while each leg is sweeping backward, propelling the body forward. Since the angle swept is always the same (front to back limit), this tells us that the legs move backward at a slower speed when the animal is moving slowly (top set), and at a higher speed when the animal is moving fast (bottom set), but always move forward with the same pattern and speed. The forward-swing phase is initiated first in the rear leg, then when that is finished the middle leg does its swing, and when that is finished the front leg does its swing. That pattern of forward swings is completely fixed in duration and sequence no matter what the animal's forward speed -- at least so it looks in the data.

The complexity of this pattern obscures a simple fact, which we see by plotting only the left rear and right rear legs:

```
Slow:
                  _           _           _
                        _           _           _

Fast:
                  _   _   _
                    _   _   _
```

This cockroach walks like three guys walking at arm's length in a column. The rear guy's left foot swings forward until it hits the heel of the guy in front, whose left leg swings until it hits the heel of the front guy whose left leg swings forward. The right side does the same thing when the rear guy's right foot swings forward and hits the right heel of the guy in front of him, and so on. If you just look at the rear guy, however, he is simply walking: left, right, left, right. The guy in front of him is walking at the same pace but slightly delayed from the steps of the guy behind: -left, -right, -left, -right.... and the front guy is walking the same way still more delayed: --left, --right, --left, --right. No matter how fast or slow the rear guy walks, the phase relationship from one guy to the next stays the same because the wave of heel-kicking propagates forward at the same speed, and the forward swing of the leg takes the same time (Bugs don't walk like human beings, so these guys are not walking "naturally." The foot snaps forward instead of mirroring the other leg. If you try walking this way, you'll find that BOTH feet are moving rearward most of the time, with left and right feet snapping forward alternately. Keep your body moving forward at a constant speed, without jerks. I suggest doing this when you're alone.).

I spent days on this, because Beer spoke at some length about the logic of moving the feet so the bug wouldn't fall down. There isn't any logic: at all times, at least two feet on each side are down, because only one foot moves at a time on each side. There's no way to fall down, no matter what the relative phases of the two sides. Most of the time, all the legs are down.

So there's something funny about Beer's model, because he says his could fall down occasionally. Furthermore, looking at the speed control circuit, p. 146, and the diagram of "phase-locked oscillators," p. 79, I can't see how the speed control circuit can reproduce the patterns above, in which the delay from each leg's forward step to the next remains constant.

For speed control, Beer uses a common LC neuron that affects all the burst-generators that constitute the stepping circuit. It appears to affect the recovery time of each P neuron in the ring of relaxation oscillators. But it is the delay between firings of these neurons that determines how fast the "metachronal wave" propagates forward from leg to leg, and this delay, according to the data for the real cockroach, is constant or very nearly so. If speed is varied as Beer seems to do it, the delay will change with speed, which it shouldn't do. In fact, although Beer shows diagrams (p. 84) in which his bug seems to produce stepping patterns like the real ones, he says in the text that the "metachronal waves" in the model break down at both low and high speeds. I think he has the wrong model for speed control.

In working out a circuit that would produce the 3-guys pattern, I ended up with essentially the same ring-of-neurons circuit Beer used, with limit sensors triggering and ending the forward steps. However, my P neurons were not one-shot oscillators as in Beer's model, but latches: the rear limit signal sets the latch, and the forward signal resets it. The circle of P neurons, with mutual inhibition, is asynchronous and does not determine the speed of walking. All it does is cause a stepping wave to propagate forward when a rear limit is reached on either side. Inhibition between the sides causes alternate stepping of paired legs, and fore-and-aft mutual inhibition assures that only one leg on each side will step at one time.

Speed control is inserted, in my version, in a layer of neurons between the P-neuron circle and the motor neurons. An LC-neuron signal simply causes the legs to move backward at a constant speed determined by the integration time of these intermediate neurons and the magnitude of the LC signal. The P-neurons and the limit sensors detect when a leg reaches a rearward limit and cause a forward-swing sequence to propagate up the appropriate side, one leg at a time, stern to prow. As each P-neuron latches, it causes the integrator driving the the corresponding leg to reset, lifting the leg at the same time. Each leg swings forward until the forward limit is sensed. This unlatches the P-neuron and enables the next one, which then goes through a similar cycle when its sensor sees its backward limit.

In the version I'm working on, the speed control circuit is not involved in stepping at all. It simply pushes the body forward at a constant speed by making the legs go backward, just as if someone were pushing the body forward from outside. The stepping circuit, consisting of the P-neurons and the limit sensors, is autonomous -- all it does is reset the leg-driving integrators to the forward position when a rearward limit is hit, also assuring that only one leg on each side moves at one time and that steps by paired legs alternate. It makes the legs look like spokes on a wheel, an image that occurred to me a few days ago before I figured out how it would work. It keeps the speed-control circuit from running out of legs to move, however fast or slow it moves them.

The P-neuron circuit looks so symmetical that I think it might work the same way going either ahead or in reverse, so that speed can be controlled in either direction, through zero. At least that's what I'm working on now. Once we have a simple reversible speed adjuster (not really a control system), we will have an output function that can be used by all the higher-level systems for moving in straight lines, turning, and reversing.

The "latch" mentioned above can be built from a single Beer neuron by using positive feedback. In the diagram below, x is an inhibitory (negative) input, and o is an excitatory one:

```
                     -------------
                    |   ------    |
                     -o| BEER |   |
       SET------------o|      |-------- OUTPUT (On or Off)
     RESET------------x|NEURON+
                        ------
```

A momentary SET input causes the output to go high, which holds the neuron ON if the gain in the feedback connnection is high enough. A sufficiently large momentary RESET signal will inhibit the neuron, turn off the feedback, and leave the neuron OFF. A slightly negative intrinsic current will keep noise from triggering the latch. So if the SET and RESET signals come from the limit sensors, the steady-state firing rate of the neuron will indicate whether a backward or forward sweep is in progress even after the limit sensor signals go away.

I haven't put these circuits into NSCK yet, but will when I get the whole locomotion thing worked out (or reinvented) to my satisfaction. I want to get reversible speed control working so that turns can be implemented by sending different speed signals to the two sides, while the stepping circuit just keeps the system from running out of legs to move.

Best,

Bill P.

Date: Tue Feb 25, 1992 11:11 am PST
Subject: Re: Quantitative vs. qualitative

(Martin Taylor 920225 13:15]
(Rick Marken 920223)

The loose canon aims at my foot, and scores a resounding hit on his own.

Rick is thoroughly confused about the problem of degrees of freedom. I am not going to try to correct his mistakes yet, unless Bill Powers fails to do so. If Bill or someone else doesn't take it up soon, then I will try once more to clarify matters. But I will point out one thing that may help Rick to more security:

>This is not true, as you can see in my hierarchical model. If set up
>properly
>(so that the perceptions are relatively orthogonal and there are
>sufficient orthogonal outputs to control them, all systems bring
>perceptions to their references; the error may not always be
>precisely zero but it remains
>within a very small percent of its possible range.
>

I am trying to say that this situation is impossible in real life. There CANNOT be sufficient orthogonal outputs to control the potential orthogonal percepts, by some orders of magnitude.

Now re-read the rest of your posting.

And mine.

In hope of mutual understanding.

Martin

PS. Why am I only now getting messages dated Feb 23?

Date: Tue Feb 25, 1992 12:26 pm PST

Subject: STUDIES ON INDIVIDUAL DISTURBANCES

++++====++++==== FROM CHUCK TUCKER 920225 ====++++====++++

I have been collecting citations of studies that involve disturbances and disruptions within interaction of two or more persons (which could be experimenter but I prefer the experimenter telling 2+ what to do) but I would like to get a list of studies that any of you use to illustrate the effect of disturbances on individual actions, e.g., elimination of feedback, "illusions", and so on. I don't care if the data have to be reinterpreted in a PCT framework; I'm more interested in the study.

Thanks,   Chuck

Date: Tue Feb 25, 1992 1:37 pm PST
Subject: Degrees of Freedom

[From Rick Marken (920225)]

I'm posting from home so I'm having some problems with the editor. That is also why todays previous post had an earlier date on it. Anyway, I will try to salve my damaged foot a bit by trying to get in synch here with Martin.

Martin Taylor (920225 13:15) says:

>Rick is thoroughly confused about the problem of degrees of freedom.
>I am trying to say that this situation is impossible in real life. There
>CANNOT be sufficient orthogonal outputs to control the potential orthogonal
>percepts, by some orders of magnitude.

Are you saying that there are more potential, simultaneously perceivable input variables than there are outputs to control them? If so, I guess that I would agree that that is probably a reasonable guess -- though I'm not completely convinced. If this is what you mean, and I accept that it is true, then what is the point? Is the point that this is evidence that behavior is the control of perception but not vise versa? If so, I don't find it very convincing. If B = f(p) where there are many more p's than Bs then what's the problem. Just have B be a function of a subset of the p's.

There are important and interesting considerations regarding "degrees of freedom" in a hierarchical control system model of behavior. I was answering your post from that perspective. One of the most interesting facts about such a model is that it will run into problems (internal conflict) if two systems are using the same outputs to control very similar perceptions (perceptual variables that are not uncorrelated -- not independent degrees of freedom). I think you were trying to make another point about degrees of freedom. I'm still not sure I understand it but I would be very pleased if you or Bill P. straightened me out on it. (But you were still wrong about the zerod reference values necessarily being sent to lower order systems from higher order systems that are successfully controlling -- keeping p=r. A perfectly controlling hierarchy of controlled perceptions will not be "zeroed out except for the highest level references". Nyaa).

Back to jury duty.

Hasta Luego   Rick

Date: Tue Feb 25, 1992 3:20 pm PST
Subject: Re: Degrees of Freedom

[Martin Taylor 920225 17:30]
(Rick Marken 920225)
>
>(But you were still wrong about the zerod reference values
>necessarily being sent to lower order systems from higher order
>systems that are successfully controlling -- keeping p=r. A
>perfectly controlling hierarchy of controlled perceptions will
>not be "zeroed out except for the highest level references". Nyaa).
>
Then I guess it is me that doesn't understand. Let's do this in concepts of one mental syllable. A control system has two incoming signals that matter, a percept (composed of some combination of potentially many different signals) and a reference (likewise composed of potentially many different signals that come from higher control systems). One is p, one is r, and they are both scalar quantities. The control system has a comparator, which compares p with r and outputs a signal e = p -r (or r - p, depending on how you like to choose your sign). The error signal is fed to an amplifier that produces an outgoing signal that serves as a reference for lower level control systems. If p = r, then e = 0, and the outgoing reference signal fed to the lower systems is likewise zero.

If a control system is in a part of the hierarchy in which all the higher ECSs that contribute to its reference signal are satisfied, its incoming reference signal is zero. It tries to set its perceptual input to zero and that is what it feeds to the higher ECSs. Also it asks the lower ECSs that provide it with perceptual input to make that input be zero.

It seems to me still that the whole stable and satisfied hierarchy will have zeros for perception and reference, except for the top ECSs in the satisfied hierarchy, whose references may be anything. But if their references are non-zero, so then must be the combination of incoming perceptual signals that combine to form their percepts. This means that the top-level non-zero references are matched by the combination function for inputs that are all zero (e.g. a sum with bias, where the bias equals the reference). Even if the perceptual function is an integrator, it seems to me that the same analysis holds.

What have I missed?

While I'm here, I might as well bore on (actually, I read an article in today's Toronto Globe and Mail that said boron in the diet can enhance intellectual function -- you get it from leafy vegetables and fruit). In his earlier post, Rick claimed that as far as he could see, an ECS with zero gain and an ECS with zero reference were both not controlling.
>
>>A control system that is not controlling is not the same as one
>>whose reference
>>signal is zero.
>

>How?
>
>>But it is the same as one whose gain is zero.
>
>Why can't it be either one?
>
If an ECS has a reference signal that is zero, and a disturbance moves the percept away from zero, then the ECS will provide an error signal that results in the percept moving towards zero. It is controlling for a perception of zero. If an ECS has zero gain, then whatever happens to the perception, the output to lower-level ECSs does not change. Nothing but external chance makes the perception move toward the reference. The ECS is not controlling.

Jury duty leaves you lots of time for thinking, in my one experience of it. I hope you are writing screeds while you wait to be called.

Martin

Date: Tue Feb 25, 1992 8:40 pm PST
Subject: Misc

[From Rick Marken (920225)]

Bill Powers (920225.0800) -- great post on cockroaches. I think I understand most of it. Are you writing your bug model in C? Why don't you return to Pascal. Must I learn a new language? It sounds like the little bugger will move quite beautifully.

Martin Taylor (920225 17:30) says:

>It seems to me still that the whole stable and satisfied hierarchy will
>have zeros for perception and reference, except for the top ECSs in the
>satisfied hierarchy, whose references may be anything. But if their
>references are non-zero, so then must be the combination of incoming
>perceptual signals that combine to form their percepts. This means
>that the top-level non-zero references are matched by the combination
>function for inputs that are all zero (e.g. a sum with bias, where
>the bias equals the reference). Even if the perceptual function is
>an integrator, it seems to me that the same analysis holds.

>What have I missed?

The fact that the output function, that transforms error into reference signal, can (and probably must) be an integrator. In my simulations, I transform error into lower level reference signal using the following simple difference equation (after, who else, Powers, in the Byte articles).

o(t+1)= o(t)+slow(gain(r-p)-o(t))

This is a leaky integration ("slow"<1 is the leak); if it were not, the integral would eventually zero if the error remained precisely zero. That did happen in some of the control systems in my spreadsheet so the absense of the leak caused problems.

If the reference input to the lower level signal really went to zero, by the way, then there would not be much to imagine in your

scheme for imagining, which I implemented in my new version
of the excel spreadsheet. In your scheme, the references to
lower level systems are fed right back up into the perceptual
functions -- weighting those reference signals using the scheme
that would have been used for perceptual signals coming from
lower levels. If the perception matched the reference then
the lower level reference would be zero -- but a (linear) weighting
of zeros is zero so the perception would not match the reference
in imagination once it matched the reference in imagination.
This line of reasoning makes me think that the leaky integration
in the output function is required.

>If an ECS has a reference signal that is zero, and a disturbance moves the
>percept away from zero, then the ECS will provide an error signal that
>results in the percept moving towards zero. It is controlling for
>a perception of zero. If an ECS has zero gain, then whatever
>happens to the perception, the output to lower-level ECSs does not
>change. Nothing but external chance makes the perception move toward
>the reference. The ECS is not controlling.

Yes, this is true. But I think of zero in the HCT model as an
actual neural current. So if e=r-p, the perceptual signal is
inhibitory and the reference is excitatory. Since neural signals
cannot go negative, if r = 0, then any p other than 0 is
inhibitory and, functionally, removes control -- there can be no
error signal and, hence, no control. But your analysis above
is correct if the r,p and e are considered algrbraic variables.

Regards   Rick

Date: Tue Feb 25, 1992 9:49 pm PST
Subject: HCT details

[From Bill Powers (920225.2230)]

Martin Taylor (920225.1730) --

I think I'd better add my voice to Rick Marken's concerning how a
hierarchical control system works. Maybe the problem here is that you're
thinking of error-correction qualitatively (either there's an error or
there isn't) instead of quantitatively. Even in a single simple control
system, the error is never "corrected." It is kept small. What keeps it
small is the fact that there is always some little error, which when
amplified becomes the drive for the active behavior that is keeping it
small. If you let the error go exactly to zero, there can be no output,
because the output is maintained by the error. When there is no output, no
action, the perceptual signal will not match the reference signal except in
a completely cooperative and disturbance-free environment. Very few
perceptions will happen to match an arbitrarily-specified reference level
without some continuing action to maintain the match.

Rick mentioned control systems with integrators in their output functions;
this could stand some elaboration. An output integrator produces an output
that keeps changing as long as there is any input to it (the error signal)
other than zero. So if the error signal is integrated, the output just
keeps increasing or decreasing until it matches whatever disturbance is
present and cancels it totally. In principle, the error signal can indeed

go to zero, because the integrator will hold the output constant until the error becomes nonzero again. But in practice this doesn't happen; the disturbance will be changing, and the system's own actions will fluctuate, effectively disturbing the perception. Every slight change in the disturbance creates some error, which causes the integrated output to rise or fall to meet the new value of the disturbance. A control system with an integrating output just acts like a slow control system. The point Rick was making was important, however. With an integrating output, a control system can, at least for a while, contain zero error, while the output is held by the integrator at a nonzero value. Thus it isn't true that when the error is truly zero in a higher-level system, the reference signals it sends to lower systems must also be zero.

In the more general case, errors never really go to zero except by accident, in passing (as a disturbance changes from plus to minus). The perceptual signal fluctuates around the value specified by the reference signal, never getting very far from the reference value (in comparison with the normal range of reference values). But the small errors that do occur are amplified and become large changes in lower-level reference signals. This amplification is the major contributor to loop gain; the higher the loop gain is, the smaller the error will be. By the same token, however, what keeps the error small is precisely that highly amplified small error, producing the changes in lower-level reference signals that (through the lower systems that receive them) counteract the changing disturbances of the perceptual signal at the higher level.

So in general, the outputs of a control system will be highly active in the process of maintaining a perception very near to a reference value. With a loop gain of 100, the effect of a disturbance on a perceptual signal will be reduced to 1 per cent of its uncontrolled effect -- but that is accomplished only because that 1 percent of the disturbance that gets through is enough to alter the system's action enough to cancel 99 per cent of the disturbance. The error in any control system at any level may be held very close to zero -- but that remaining error is what accounts for most of behavior.

Now what about control systems at higher levels, where the variables do tend to be qualitative and jump from one value to another (class membership, for example). Suppose a man says "I've been rich and I've been poor, and believe me, rich is better." So he sets his reference level to "rich." Now it seems that either there will be an error (he's poor, meaning not rich) or no error (he's rich, meaning not poor). In terms of logic and discrete variables, this control system either experiences an error or it doesn't. But that's just a way of talking; that may be how language seems to work but it's not how perception works.

Suppose this man earned $500,000 after taxes in 1990, and managed to save $400,000 of it by frugal living. He starts 1991 with $400,000. Is he rich now? Well, he's richer than he was in 1989, but some of his friends made $750,000, so actually he doesn't feel as rich as he might have if they had only made $200,000. Somehow this either-or variable called "richness" has developed some elasticity -- a scale of continuous variation is creeping into the perceptual picture, and it's not suppose to do that for categorical variables. The reason it creeps in is that "richness" is really a continuous PERCEPTUAL variable, but we don't have as many symbols for degrees of richness as we can perceive degrees of richness. In fact, people get very lazy when they talk about perceptions -- it's much easier to use either-or terms that make it SEEM as though categories are neatly discrete and stay where they are put.

If we use symbol-manipulations as a guide to behavior, we may set up control systems that seem discrete. The problem is that they don't control very well unless we use quantitative symbol systems like mathematics. If the only reference values you have are "present" and "not present," you can't even perceive an error until the underlying perceptions have crossed the boundary of a category. So disturbances that actually alter perceptions almost to the boundary aren't recognized and no action is taken.

Furthermore, actions themselves will then be driven by on-off error signals: either you act or you don't act. So it's a certainty that you will always underrespond to disturbances or overrespond -- you'll seldom hit the amount of response required by a specific disturbance on the nose. Control will simply be very poor: the loop gain, in fact, will be about 1, so you'll cut the effects of errors in half. That's all.

This is what I have against made-up control systems like decision theory. You decide on a course of action using the best reasoning processes available, commit yourself to it, and the next day there's a disturbance that calls for a slightly different decision that doesn't get made because method can't resolve small enough differences. Qualitative either-or thinking always runs into this problem: the universe changes by degrees, but qualitative variables can change only in big jumps. Whatever they represent, therefore, can be controlled only to the extent of putting up a feeble token resistance to disturbances.

So while qualitative errors may seem to be zero at times, they aren't really zero but are only treated as if they were zero. They may result in no action at a lower level, but that's not because there's no need for action. By the time you decide that the "bull market" has changed to a "bear market," it's too late. The guys who made all the money were monitoring continuous variables, not categories.

All of this underlines a basic aspect of HCT. Reference signals do not cause action. Neither do perceptual signals. Action reflects error signals, the difference. And error signals can drive wildly fluctuating actions at lower levels which see to it that the absolute amount of higher-level error remains negligible, on the time scale appropriate to the control system.

One other point. You've been saying that there are far more potentially controllable perceptual variables than there are degrees of freedom of action. This is true. There is, in fact, an INFINITY of POTENTIALLY controllable variables, because a new potentially controllable variable comes into existence every time a perceptual function reorganizes. How many relationships can you see between your nose and objects in the environment around you? How many KINDS of relationships? There's no limit -- you can go on noticing new and potentially controllable relationships for the rest of your life, if someone feeds you. Those relationships aren't Out There waiting to be discovered: you make them up. As soon as you do, you can try to affect them by some action. If the action affects the relationship, you can get organized to control the relationship. If controlling any particular relationship leaves you better off, you'll continue controlling it, and find the particular reference state for it that's best of all for you. If nothing good happens, you'll abandon it.

And what about the effects of your actions BEFORE you notice a particular relationship? They must have been affecting it then if they affect it now. But so what? The only ones that can make a difference in your actions are those you perceive -- the rest may as well not exist.

I think the problem here may be that you're taking the external view of, for example, MY controlled variables. You can see lots of things I might be controlling for, and you can see that my mere 800 muscles couldn't possible handle them all at once. But if you could look at matters from in here, you'd see that I have selected for control, in fact have created, just those variables that make a difference to me. All the rest are known to me only if they happen to affect my intrinsic state, and then I know of them only by feeling bad for reasons I don't consciously understand.

Another factor that needs to be considered is context. I can in fact control far more than 800 degrees of freedom of perception, as long as I don't try to control them all at the same time. Whatever output degrees of freedom I used in eating dinner are now free to be used in typing; I can't do both at once, but I can schedule them. In fact if I'm willing to give up a little precision of control, I can even interleave them, grabbing a bite and typing a sentence, and so on. Also, if I'm clever, I can pick actions that satisfy multiple goals, improving my touch-typing (ha!) at the same time I'm writing something supposedly useful. That is a way of reducing the degrees of freedom, of course, but we do that all the time.

The main thing, however, is to realize that at any given moment, certain perceptions are under active control, and at that time there is no degrees-of-freedom problem at all. The perceptions exist; they depend on the world in a certain way; they are affected by our actions on that world. For any given perception, all that matters is that the available actions be able to alter the perception to make it more like the reference signal. In the course of doing this, the external world may change in a thousand ways that are irrelevant to this one perception; that doesn't matter to the control system. The world is left to change as it will in all those dimensions that are not being perceived or controlled. There's no need to make all the variables in the external world behave just so. All that's required is that the perceptual signal that results in one particular system be affected in the right direction.

Last remark. When I was 18 years old I was in the Navy learning to be an Electronic Technician's Mate (we joked about the time when we finally had to meet the Electronic Technician). One of the things an ETM had to learn was how to troubleshoot negative feedback systems, which were found everywhere from gun-control systems to radio receivers. Our instructor introduced us to the control equations (which were to mean a little more to me nine years later) by showing how raising the gain of a negative feedback amplifier improved the linearity and frequency response of the amplifier. At one point he showed us that when the loop gain was made high enough, the characteristics of the vacuum tubes and audio transformers became less and less important, until in the limit the error signals were all zero and the amplifier properties cancelled out of the equations. So, he said, if the vacuum tube gain no longer makes any difference and has canceled out of the equations, why can't we just pull out the vacuum tubes and let the feedback do the work?

There was a considerable pause before the laughter, because we had all worked through the equations and seen the answer -- and believed it. It took a moment to realize that without those vacuum tubes there wouldn't be any loop gain, even though their exact properties no longer mattered much.

From this we learned something -- to be suspicious of handy rules of thumb, and always to remember how we got to them. That was a rather fine instructor. He always told us not to just go through the motions, but to understand what we were doing, what it meant. I got more out of that Navy course than I did out of most college courses later on.

So, passing on the good word, I'm saying that to understand HCT, one can't afford to settle for just the rules of thumb, the approximations that allow us to talk in handy qualitative terms about a system that is quantitative to its core. One has to grasp the interplay between error, output, action, perception, reference signal, and error again. One has to understand that "error correction," "goal direction," "stabilization," "resistance to disturbance," and all such easy shorthand phrases are just verbal approximations to a process that's not as simple as that. That's why I keep on so tiresomely about trying the tracking experiments, being a control system, pondering the details of what goes on, so you'll see directly what is happening and not reduce the process to slogans -- even slogans like "control of perception." An abstract understanding of control processes isn't enough to create that deep understanding that comes from wordless experience of the dynamic relationships and your own personal part in them.

Best to all,

Bill P.

Date: Tue Feb 25, 1992 10:11 pm PST
Subject: Action theory

[From Rick Marken (920226)]

This jury duty stuff does orivide a lot of time for fooling around. I should be trying to write my book -- but it's just too trying. The net is much more fun anyway. So here is a belated attempt to answer Gary Cziko's student.

Gary Cziko (920222.1730) says:

>I student of mine recently showed me an edited volume:

>Frese, Michael., & Sabini, John. (Eds.). (1985). _Goal directed behavior:
>The concept of action in psychology_. Hillsdale, NJ: Erlbaum.

>It contains 23 chapters, divided between Americans and Germans
>including two names I recognize, Gallistel and Neisser.

You're not going to believe this, but there is a guy in my jury group named Gallistal. He runs around the waiting room with a portable Mac and an introductory Psych text. Could it be the same guy? If so, what is his point of view? Should I ask him anything about "action theory". This could be too weird.

>In the editors' introduction, they offer a definition of action theory:

>"Action theory begins with a conception of human behavior: that it is
>directed toward the accomplishment of goals, that it is directed by plans,
>that those plans are hierarchically arranged, and that feedback, from the
>environment articulates with plans in the guidance of action." (p. xxiii)

>I can see how on first glance such a theory may seem to be very similar to
>perceptual control theory

Well, it's not really a theory. It's more like a sentence. I've heard people say things that were breathlessly close to control theory many times before (indeed, most reviewers of my stuff say that they already know that). But when the big crunchola comes (the idea that behavior is controlled perceptual variables) there is silence, dismay, hostility-- all those things that enrich the life of the lonely PCTer.

> but I hope I am sophisticated enough by now to
>see where action theory (or at least this conceptualization of it) is
>problematic and how it is NOT the same as PCT.

You bet you are!!

>But I may need some help in explaining this to my student.

I doubt it.

>Any of you PCT old-timers (or anybody else) care to help out?--Gary

I'm in.

>P.S. Does the distinction between action and behavior made recently by Greg
>Williams apply to the action of "action theory" as well? (I could do with
>some more explanation of this disctinction, Greg.)

Maybe. I was interpreting Greg's "actions" as including both effects on controlled variables (the "real" PCT definition) as well as interesting (to an observer) side effects. I have no idea what the "action" in "action theory" means. But as I re-read the definition above, it sounds like someone is going out of his/her way to describe behavior as "caused output" while admitting that it is "goal oriented" and may need some buffing up with some of that good ol' feedback stuff. Just ask if they understand the concept of a controlled variable. See what they mean by "goal". How are these goals achieved in a disturbance prone environment? You know the litany. Best of all. ask them "what the **** is behavior?" (in those words). See what they say. That ought to nail it.

Love Rick

Date: Wed Feb 26, 1992 8:40 am PST
Subject: still more bridges

From Greg Williams (920225)

(Regarding Bill Powers (920222.1000))

Sorry I'm getting a bit behind -- busy typesetting LCS II.

I'm concerned about THIS kind of "detail" in PCT models: Suppose you model a person's tracking controller as a first-order differential equation, a lag. Suppose further that it does a good job of predicting empirical tracking data when disturbances aren't very trying. I claim that such prediction doesn't tell you enough to decide whether the first-order model is reasonable instead of, say, a second-order model. The experimental data isn't critical enough with regard to choosing between candidate model structures. But suppose you suddenly change the cursor position (a step change resulting in a transient "response"). Does the subject overshoot? If so, then you know you that a first-order model isn't correct. In other words, my concern is that psychologists will not be very impressed by accurate predictions (after all, an infinite number of control models could give good predictions when control is good; choosing among them via Occam's razor isn't a very convincing procedure!), but WILL be impressed by data allowing some PCT models to be favored over others (allowing PCT models to be favored over nonPCT models is a separate issue, and, as you imply, maybe a prior one -- my enthusiasm about Rick's plans was meant to address that issue).

Again, I am not suggesting that attempting super-precision in predictions made by PCT models will sway nonPCT psychologists. Rather, I am suggesting critical experiments to flesh out the details of PCT models at the "structural" level (i.e., is it P- or PID-control) rather than at the "parametric" level (i.e., is the loop gain 20 or 30). This is what nonPCT psychologists are shooting for with their nonPCT models, as far as I can tell from (some of) their papers. It's time for PCT modelers to start outgunning them!

>But I say flatly that there are no such models other than the CT model, and
>the CT model shows that when you can predict actions accurately, you're
>just predicting disturbances. Show me that I'm wrong. Whose definition of
>"accurately" are you using? Have you got an example of an accurate non-CT
>prediction of actions?

I suspect that you would be somewhat surprised by some (not many!) of the nonPCT models in, for example, BIOLOGICAL CYBERNETICS, with respect to their sophistication regarding predictions of various organismic actions.

When control is good, actions reflect disturbances accurately. When control is poor (hard or going toward good control in transients), the dynamics reflect a combination of disturbances and the organization of the organism's control structure. The poor-control situation is the ONLY situation where you can begin to identify the details of the control structure -- in the good-control situation, you (almost entirely) see just the disturbance-mirroring.

Best,

Greg

Date: Wed Feb 26, 1992 12:38 pm PST
Subject: Beer on Powers

[from Gary Cziko 920226.1340]

I took the liberty of sending Randy Beer excerpts of two of Bill Powers's posts on the Beer bug (Powers 920222.1000; 920225.0800). I also told him that I had another half a megabyte or so or posts from CSGnet relevant to his work that I would send him if he was interested. Below is his reply.

Since he didn't request the other posts, I will not send these to him. But I will send him information on Powers's two books.

If anyone wishes to communicate directly to Beer, his address is <beer@cthulhu.ces.cwru.edu>.--Gary

==
[from Randy Beer 920226]

Gary,

Thank you for continuing to send me excerpts from the discussion on CSGnet. Unfortunately, with my teaching responsibilities this semester and a number of other committments I've made, I'm afraid that I don't really have the time to actively participate in these very interesting discussions on a regular basis. However, since you've taken the time to send me these excerpts, let me try to reply to at least some of the issues that have been raised.

It is certainly the case that the control systems of the artificial insect could be described in a more abstract form as Powers seems to be trying to do. However, one of the main goals of my work was to demonstrate how the various behaviors could be implemented in neurobiologically plausible ways. To take just one example, Powers states that "The `overeating' phenomena, by the way, can be modeled as a delay between the ingestion of food and the subsequent rise in energy level. No `overeating' neuron would be needed." This statement may be perfectly true, BUT THAT'S NOT THE WAY BIOLOGY DOES IT, at least not in the animals whose feeding circuitry has been studied. The design of the consummatory circuit was directly based upon work on the neural basis of feeding in the marine mollusc _Aplysia_. That animal has an identified neuron (the C2 neuron; for a recent review of this work, see Chiel, Weiss and Kupfermann, "Multiple roles of a histaminergic afferent neuron in the feeding behavior of Aplysia", _Trends in Neuronsciences_ 13(6):223-227) that appears to play a role very similar to that of the feeding arousal neuron in the artificial insect.

Aside from the fact that, once again, there is no behavioral evidence that insects walk faster the hungrier they get or the greater their distance from a patch of food, I was a bit confused by the specifics of Powers proposal for using the total odor strength. How does one choose the reference odor signal which is supposed to initiate ingestion? Since the odor strength of a patch of food is proportional to its size and decays as the square of the distance from the patch, any fixed reference signal that would work for one patch would cause the insect to begin biting before it had actually reached a larger patch. Making this reference signal dependent on the level of hunger only complicates this problem, because the hungrier an insect gets, the harder it may be to find a patch large enough to satisfy this reference signal, causing it to walk right over smaller patches that it might encounter! The reason that I used a multimodality sensory trigger (tactile and chemical sensors) for injestion was precisely to deal with problems such as these.

There also seems to be a misconception regarding the interaction between edge-following and the appetitive phase of feeding. The insect will not "...ignore [food] to go edge-following". As explained in Chapter 8, edge-following takes precedence over the appetitive behavior only if an obstacle lies between the insect and a percieved source of food. Note also that finding its way around obstacles to food is not the only reason the artificial insect possesses an edge-following behavior. In fact, cockroaches are thigmotactic, and spend most of their time within antennal contact of an edge, even in the complete absence of food.

Turning to the locomotion controller, Powers correctly observes that normal gaits consist of overlapping metachronal waves on each side of the body (this is a direct consequence of Donald Wilson's analysis of insect gaits, which is discussed on p. 72-73 of my book). However, it is not true that only one leg swings at a time. At higher speeds of walking, the metachronal waves on a given side of the body begin to overlap, so that the rear leg begins a new metachronal wave before the previous one has been completed (this is obvious for the tripod gait, but is also true for most of the other gaits; see Figure 4.5). All of these gaits have the important (to the insect!) property that static stability is continuously maintained.

Other miscellaneous points on the locomotion circuit: (1) The delay between subsequent swings in a metachronal wave is set by the membrane time constants of the pacemaker neurons and is therefore largely independent of walking speed. (2) The tonic command signal affects only the burst period of the pacemaker neurons (again, in a way consistent with biological pacemaker neurons). (3) The insect occasionally falls down because, as in most systems of entrained oscillators, phase-locking is not globally stable, but only locally so. Sufficiently large perturbations (such as large instantaneous changes in speed) can upset the entrainment underlying the metachronal wave. Also related to this is the whole issue of turning, which upsets the clean stereotypic patterns observed in straightline locomotion because stancing legs on the outside of a turn have to swing through a larger angle than those on the inside of the turn. (4) I certainly won't disagree with the statement that I have "the wrong model for speed control." Insects utilize hundreds to thousands of nerve cells in the control of locomotion, while my model has only 37. It is clear that my model circuit is oversimplified in many ways. However, it is consistent with most of the currently available facts, which is the best that I can hope for at this point.

I encourage Powers to pursue his alternate gait generation model, and will be interested to see how it turns out. At first glance, it sounds similar to the gait controller that Rod Brooks used in Genghis (his first generation walking robot). It is essentially a peripheral or "reflex" oscillator which uses sensory signals to trigger the transitions between stance and swing. Again, the point of my work is not that such controllers cannot be designed, but that such approaches are not generally utilized by animals. Central pattern generators with (often extensive) sensory modulation seem to be the order of the day in neurobiology. The obvious drawback to a purely peripheral oscillator is that sensory damage can shut down the entire system, while animals (and my locomotion circuit) can still generate the basic rhythmic motor patterns in the absence of sensory information. Of course, this is not to downplay the very significant role that sensory feedback usually plays in reinforcing and fine-tuning the overall pattern.

More generally, let me say that there are plenty of shortcomings to the work described in my book. The neurobiological connections are sometimes purely speculative due to lack of data, they are often more ad hoc than I would like, and many simplifications have been made in order to put together a complete (albeit simple) "animal". But most of my decisions have been motivated by the desire to remain as consistent as possible with the relevant (neuro)biological details. Of course one can debate whether THAT is the right thing to do, but that is a different discussion. I'm also afraid that my appreciation of some of Powers' points is hampered by my lack of familiarity with PCT and HCT. It sounds as if PCT and HCT might be quite relevant to my research interests and I would certainly enjoy learning more about it when time permits. Is Powers' book widely available? Our library does not seem to have it.

Sincerely,
Randy Beer

Gary A. Cziko   Telephone: (217) 333-4382

Date: Wed Feb 26, 1992 4:46 pm PST
Subject: Re: HCT details

[Martin Taylor 920226]
(Rick Marken 920225 and Bill Powers 920225.2230)

I find myself in a peculiar mental state, which I can imagine to be rather like Bill's when trying to deal with classical psychologists: I've got this great way of helping them to see their work in a new light, and why won't they accept it? I feel that I am trying to provide substantial strengthening to the PCT model by showing how it follows from first principles, and the main protagonists react as if I were trying to dissolve their fragile edifice. Non comprendo.

Between Rick and Bill, three issues seem to have separated themselves out of my original "degrees of freedom" posting:

1. Rick and Bill both argue that if the error signal of an ECS is integrated to provide the reference signal it sends to lower ECSs, then the lower references in a stably controlled hierarchy need not approach zero.
   Point 1 is really unimportant in the general scheme of things, but it is an interesting byproduct, which may help in teasing out notions of conscious and unconscious perception. It could become important later, so I leave it in the discussion.

2. Neither Bill nor Rick accept the point that the degrees of freedom argument implies that there are uncontrolled but potentially controllable percepts.
   Point 2 is crucial, and we really should come to an agreement here.

3. Bill brings up a new point that the boundaries of categorical decisions are unstable and context-dependent, and therefor that the continuous underlying percepts form the subject of control, not the decisions based on the categorical boundaries.
   I don't know where this came from, but it is interesting, and actually conforms to a description of cognition and the relation between perception and logic that I had developed a couple of years before discovering PCT.

I will address Point 1 in this posting, leaving Points 2 and 3 to later posting(s).

On 1:
     Both Rick and Bill describe integration that can be seen as a 6 dB per octave low-pass filter. Such a filter is (in my mind) the mirror or partner of the "slowing factor", and is a necessary aspect of the differences in

bandwidh between higher and lower ECSs. Both promote stability, in part by reducing aliasing related to discrete time-sampling, and more generally by ensuring that reference signals are not sent from higher levels at rates that could cause phase-shifts that could generate oscillations at the lower levels. I use much the same concept in my dynamic neural networks. But I fail to see how this changes in the slightest my argument.

Bill adds that there are disturbances in the environment that shift the percept away from the reference, and that small errors at higher levels translate into big activity at lower levels. But the same argument should hold in reverse--good control at lower levels should result in microscopic errors at higher levels. Furthermore, I postulated the same conditions Bill often does--slowly varying disturbances that remain under tight control. So again I can't see the objection to my original analysis.

So far as I can see, unless there is a bias (and Rick proposed one in the thresholding of a neural current at zero), all the percepts in a stably controlling portion of a hierarchic control system will stay very close to zero. I lay awake many hours last night trying to see another stable solution to this organization, and failed. That doesn't mean there isn't one, but I can't see it at the moment. Bill and Rick must show that my condition is actually not a solution at all. Integration that is simply low-pass filtering does not do it. Non-linearities in the comparator responses might do it, but I have my doubts. (Zero neural currents do not necessarily correspond to zero percepts; most neurons have a non-zero resting rate of firing, which reduces if they are inhibited, increases if they are excited).

Let's try this one for size. I have to go now, in rather a hurry, or I would continue. Points 2 and 3 later.

Martin

Date: Wed Feb 26, 1992 6:56 pm PST
Subject: learning grammars

Here's how learning verb lexical entries *might* go. First, we suppose that children build up some sort of table of correlations between (representations of) utterances, and (representations of) the situations that these utterances seemed to be about (using pointing & direction of gaze as evidence concerning aboutness. E.g. `Do you want a cookie?' correlated with a memory of daddy holding out a cookie. From this, a table of correlations between features of utterances and features of situations is worked out (e.g. meanings, of a sort). This task is made easier by (a) restricting the relevant features of utterances to continuous sequences of phonemes (or maybe even syllables, at the earlier stages) (b) the presumed fact of a rich, preexisting scheme for classifying the situations, in terms of individuals, actions, etc. E.g. when Owen sees Jimmy break a glass, it is already classified as something Susan does to something.

So when we come to an actual utterance such as:

Jimmy broke the glass

The child already knows more or less what `Jimmy', `glass' and `broke' mean. Now at this point different actual linguistic theories will procede differently, as befitting their difference conception of what grammars are like. In LFG it is assumed that there are PS rules and grammatical relations, and that the major kinds of ingredients in situations each have a standard form of expression in the syntax, even though it doesn't have to be used. Furthermore the `Xbar theory' tells you that for each grammatical category of word (`part of speech'; `lexical category' there are various different kinds of phrasal nodes (what kinds depends on what flavor of Xbar theory). One popular idea is that there are two kinds of phrasal nodes, a `maximal one', and an intermediate one.

One of the basic perceptual categories is THING (including PERSON as a subcagetory), whose associated syntactic category is noun, N, so Jimmy and glass are N's. Due to Xbar theory each of them sits under an intermediate nominal category (often called NOM, and notated N', or N-with-one-bar-on-top), and a maximal nominal category usually called NP (or N'', or N-with-two-bars-on top), so this gives us:

```
      NP                     NP
      N                      N
    Jimmy   broke     the  glass
```

Now `broke' is also known to be associated with an EVENT, whose associated grammatical category is verb (V), so we will have intermediate and maximal verbal phrasal nodes as well (VP and S). The NP's also have to get grammatical relations in order to be interpreted as arguments, that is represented as playing roles in the breakage event. Furthermore Agent- and Patient-like grammatical relations are each expected to be associated with a different grammatical relation, which we label SUBJ and OBJ, respectively (note carefully that all the theory has to say is that various kinds of ingredients of and ways of participating in situations are normally associated with some syntactic category - the actual labels for these categories can be seen as chosen arbitrarily).

So now we have motivation for putting SUBJ on the first NP and OBJ on the second. For a final flourish, X-bar theory doesn't like to put `heads' of phrases in the centers of those phrases, but doesn't mind have one argument outside of the intermediate phrase (the `external argument' in GB parlance), and is furthermore biassed to favor putting Agents externally. So the tree structure for our sentence becomes:

```
                S
 NP:SUBJ              VP
  NOM         V            NP:OBJ
   N        broke            NOM
 Jimmy                        N
                           the glass
```

The PS rules then need to be abstracted off the structures (perhaps on the basis of generalizing in some fairly narrow way over the structures that have actually been posited, so that if NP VP is all one sees under S, NP VP is what is posited), but the form of the lexical entry of the verb, PRED `Break(SUBJ,OBJ)' is fixed by the fact that that's the only one whereby the structure will actually mean what it has been guessed to mean.

It should be clear that the details and difficulty of these steps will be highly contingent on the details of the theory of grammatical structures and rules: if these are wrong, it will all be much more difficult than it has to be. That's one reason why linguists have not worried very much about mechanisms for language-learning (but c.f. the Pinker book I've mentioned before, and a recent book by Andrew Radford, Syntactic Theory and the Acquisition of English Syntax, 1990, on acquisition and Government-Binding theory), believing that it is more important to get good methods for grammatical description first.

Avery.Andrews@anu.edu.au

Date: Wed Feb 26, 1992 9:41 pm PST
Subject: Reply to Beer

[From Bill Powers (920226.2100)]

Dr. Randy Beer (920226) --

Good to hear from you, if only indirectly. You're probably right not to get involved directly with CSGnet if you want to get anything else done, but perhaps with Gary Cziko acting as filter a more modest level of discussion could be maintained. Despite all our enthusiastic second-guessing of your model, it's evoked a lot of interest and some authoritative comments from the source would be of interest to me, and I'm sure all the others.

A few comments:
You say

>It is certainly the case that the control systems of the artificial
>insect could be described in a more abstract form as Powers seems to
>be trying to do. However, one of the main goals of my work was to
>demonstrate how the various behaviors could be implemented in
>neurobiologically plausible ways.

I don't see any incompatibility, as I am basically a hardware person myself, with software second and abstractions last. Actually, my use of terms like "input function" and "comparator" are no more abstract than other terms like "sensory nerve" or "motor nerve;" it's simply a way of characterizing the function that's performed by an active component. A comparator basically emits an output signal that's a measure of the difference between two input signals; in electronics this function is carried out by a "differential amplifier," and in your model by a neuron that receives a positive and a negative input current. Of course such a device is a comparator only when it appears in a particular place in a functional diagram.

I, too, am interested in "the way biology does it." Many features of HCT were derived from a study of neurological literature (within my capacities), and I've tried to avoid postulating arrangements that would be neurologically impossible or contrary to known organizations. I think you'll agree, though, that what a circuit does depends even more on the detailed parameters than on the physical connections. When neither the parameters nor the connections are well enough known to force a particular design on us, we have to fall back on inventing designs and figuring out what functions need to be performed. Most of the HCT model is concerned with defining plausible functions -- the neural details will have to wait a while.

I had better explain my observations on the "appetitive" system a little more fully. I did not mean to cast doubt on the existence of the "arousal" neuron, but only to suggest a different interpretation of its function. It is, in fact, organized as a comparator. An overview of the whole appetitive-consummatory system might look like this:

```
                            Reference input
                                   |
                            -      | +
     Energy level -> Sensor -> comparator -> food-seeking/ingesting
          ^                 (Arousal neuron)         |
          |                                          |
          |        Feedback effects                  |
           -------------------------<---------------
```

The "overeating" to which I referred was the effect of eating more than is actually needed to bring the energy level up to specs (I believe you mentioned this in the book). If there's a delay between ingesting and an increment in the sensed energy level, this system will ingest more than is actually needed to correct the energy deficit. Note, by the way, the explicit completion of the feedback loop above. I think it's important to include feedback effects in all such diagrams, even when they involve the environment.

In your model the reference input is implicit in the "spontaneously active feeding arousal neuron" (p. 128). You have elected to make this spontaneous level of activity intrinsic to the neuron, but there's nothing in principle to say it couldn't be determined by an actual input current rather than an intrinsic current. If in fact there are no positive inputs to this neuron in the real animal, I would accept a fixed reference signal inherent in the neuron; otherwise I would be looking for a way in which other systems could set the base level of spontaneous activity, through an excitatory input signal, as shown above.

Whatever the case, this neuron puts out a signal that is maximum when the sensed energy level is zero and that falls to zero as the amount of inhibition from the sensor approaches the amount of intrinsic (or external) excitation in the arousal neuron or comparator.

My suggestion about using this signal for (among other purposes) setting a reference level for total odor input was an attempt to make a control system that would bring the animal to a food patch and stop it there, without the need for inhibition or deactivation of the food-seeking system. As you have observed, the odor strength will be an inverse-square function of distance between bug and food. What I visualize is roughly this:

```
                 (Ref?)
                   |   (Err signal)
Energy -------> Arousal ------------>---------------------> biting logic
Sensor          Neuron                  |                 and large loop
  |                                     |
  |                                     |(functions here as
  (effects of biting)                   | reference signal)
                                        |
                                      + v     (Err signal)
                                     Comparator --------->Walking speed
                                      Neuron                  |
                                      - ^                     |
                                        |                     |
                                    odor|sensors        distance to food
```

```
                                    |                    |
                                     -Inverse square law<-
```

The arousal neuron is a comparator in a higher-level system concerned with control of energy level. Its output becomes (among its other functions) the reference signal entering another comparator in a lower-level system. This second comparator receives a sensory signal representing total odor intensity, and emits a signal that is the difference between arousal output signal (now odor reference signal) and odor sensor input signal. That error is maximum when the bug is far enough from the food so the inverse-square attenuation makes the odor intensity nearly zero. The error signal will remain essentially at maximum until the bug is close enough to the food that the odor intensity rises significantly along the nonlinear inverse-square curve. As the odor intensity rises, the inhibitory sensory signal going into the second comparator rises, and the drive to the walking speed generator begins to decline. We can assume that directional control works as per your design.

By adjusting thresholds and sensitivities, you can make the odor error signal begin its decline when the bug is within any radius of the food. Within that radius the bug will simply slow down and stop when the odor intensity matches the reference signal set by the output of the arousal neuron. Of course you could add proximal chemical sensors to this loop to raise the total "odor/taste" signal rapidly when food is physically encountered. One nice side effect will then be that as food is eaten, the bug will move again until more food is encountered in the patch.

The lower the energy level, the higher the reference signal going to the odor level control system, and the faster the bug will move (of course a saturation level can be set to limit maximum speed of movement). A minimum speed can be set such that with an energy level matching or exceeding the reference level, the bug will still move slowly -- or one could add other higher-level systems concerned with control of other important inputs, also making use of the speed-control system.

If bugs don't walk faster the hungrier they get, all you need to do is tailor the energy-level comparator accordingly: it puts out maximum signal until the energy level is nearly at the reference level. The odor control system doesn't affect energy level; only biting will raise that to the reference level. But the bug will still stop walking when the odor signal rises to the reference level set by the output of the arousal neuron.

As you can see, this proposal doesn't imply that bugs walk faster when farther from the food -- only that they walk slower when nearly at the food, and stop when they get there. When far from food, the inverse-square law means that speed will remain nearly constant until the near vicinity of the food is reached; then it will decrease. You can adjust the parameters to make the slowdown occur at the same distance it occurs with the real animals. I'm sure they do slow down and stop when they get to the food! The difference between my proposal and yours is that mine doesn't require deactivation of the approach system by superordinate logic. It can work all the time. If the odor patch moves while the bug is eating, the bug will simply follow it without any special instruction to do so.

I will try to implement this design using Greg Williams' "NSCK" program. If you beat me to it, fine by me -- just let me know how it works.

I think we can leave other details of your remarks for later.

Best regards,

Bill Powers
P.S. I'm posting this on CSGnet and sending to you directly. You can reply through the net, CSG-L, as your comments will be of general interest.

Also, note my odd address: I'm piggybacking until I get my own logon.
WTP.

Date: Thu Feb 27, 1992 3:29 am PST
Subject: Bill's bug

From Pat & Greg Williams (920227)

Just a quick note of support for Bill's PCT-remodeling of Dr. Beer's bug. Exactly what we've been trying to encourage (not enough time to try it ourselves, unfortunately). We look forward to seeing Bill's bug "haul cerci."

Gary, your timing on sending Bill's posts to Dr. Beer was perfect. Excellent discussion of the issues!

Now if we can only get a couple of behaviorist bugs (SR and operant, of course) for the great "Bug-Off"....

Best,

Pat & Greg

Date: Thu Feb 27, 1992 5:17 am PST
Subject: csg/biblio.pct

Following your discussion about PCT and Beer's Bug I tried to get both books.

In november 1991 I ordered the Beer-book via Lonodon (there is an european distributor of Academic Press) ... I'm still waiting!

The Power's Boook seemed to be simpler to get; I followed the advices from file biblio.pct (biome-server) for international order:

Powers, W.T. 1973. _Behavior: The control of perception_. Chicago: Aldine.

The book that started it all--seminal explanation of control theory. This is why CSGnet exists.

ISBN 0-202-25113-6, 1973, 296 pages, hard cover, $38.95 + 3.50 (shipping & handling; NY residents please add sales tax) MC, VISA, money order.

Ordering Address: Aldine de Gruyter, 200 Saw Mill River Rd., Hawthorne, NY 10532
Phone: 914-747-0110 Fax: 914-747-1326
International Order: (Berlin) fax 011-49-026005251
(About 150 copies still available)

I tried to get it from Berlin (international order) but my order was cancelled today: there is not one copy available.

What remains?? I'm frustrated!

Wolfgang Zocher
Regionales Rechenzentrum
Universitaet Hannover

Date: Thu Feb 27, 1992 8:53 am PST
Subject: Language development; ordering BCP

[From Bill Powers (920227.0900)]

Avery Andrews (920226) --

Your post suggests much about links between HCT and language learning. One phase seems to be naming of low-level perceptions: sensations through events. So naming itself -- the substitution of a word for a non-word perception -- appears to occur while the hierarchy is still primitive. This raises questions I can't answer, but perhaps people in developmental linguistics might know some answers. Or perhaps someone with a CT orientation needs to do some experiments!

One question is whether naming begins concurrently with the ability to control general perceptions at the various lower levels. I suspect it doesn't. An infant has time to develop skills up to some level before even basic naming starts. Is that level as high as relationships? If it is, the suggestion is that naming is the first sign of acquisition of a level of control above relationships -- categories. If not, it suggests that naming is not a process of categorizing but something simpler, perhaps a level omitted from the HCT model. I don't see any theoretical basis for deciding. Someone just has to apply the Test to determine what levels of variables a child can control, and compare that evidence with the growth of speech.

Within the boundaries of linguistics, we could ask what level of speech control exists at various ages. Imitating sounds is clearly control of sensations; when the sounds take on specifically linguistic character, we might think configuration control and transition control are appearing. The utterance of word-like sounds suggests events. Of course the skills at making sound-sensations, word-sounds, and word-events spread sideways with practice, so the repertoire becomes wider and the reproduction of heard word and phrase events more precise.

The relationship level would seem to represent the commencement of using one thing in relation to another: a hammer for banging on things? A word to indicate an object or event? Is this where naming of specific perceptual items begins? If so, basic naming doesn't involve categories -- the child doesn't mean "a" dog by the word "dog", but THAT SPECIFIC dog, that perception. The substitution process would then be a relationship-control-like process.

Alternatively, naming could require the category level, with the most elementary category being composed of a spoken word and one specific perception. As the category level develops, a word can be attached to categories with more than one nonverbal member. Adults use the same word, dog, for a number of different configurations, so all those configurations become members of the same category. They even show pictures or toys and say "dog." So the category gradually drops variable and inconsistent lower-level perceptions and retains only the features always present when "dog" is used. At this stage, use of multiple words (Ball Jump Dog) is not sequentially significant; the sequence is simply forced by the lineal nature of speech, and the order in which referents are noticed.

The next thing to enter, according to HCT, would be sequence as a specific controlled variable. Now Cat Bite Dog is different from Dog Bite Cat. Is this the stage at which nonverbal control also shows evidence that sequentiality has become a controlled variable?

Finally (for now), we get to the program level where rules enter into control. That is, sequences are formed to satisfy the requirements of logical, grammatical, or other kinds of rules: if this, then that, else the other. At this point you get complete sentences that fit general forms such as NP: SUBJ VP. At this stage, we should also see children beginning to do other things according to logical rules: If the cap is on the bottle, untwist it before pouring the syrup on the baby, otherwise just pour (no principles yet).

This progression seems to fit, more or less, the stages you present. The question I'm raising is whether we could think of a research project in which the general capacities for control are investigated and tied to the specific stages of language growth in the terms you've suggested. I'm sure nobody has yet done a developmental study using the Test and specifically looking for the growth of skills (Frans and Hedwig Plooij have done studies of the development of levels of control in chimpanzees, and are now working with human infants, but haven't, as far as I know, gone into language per se). Do you think there might be indirect evidence in existing developmental studies? Or would it be too unlikely that the specific levels of HCT would have been used, or that actual testing for control would have been done?

A study like this could go far toward answering the question of whether language is a specific hierarchy of skills, or simply a specialized use of general levels that pertain to all kinds of behavior. Even if such studies have been done, the Test for the Controlled Variable would introduce a new tool and a new conception of just what skill-development means. Another candidate for the the list of research projects?

Wolfgang Zocher (920227) --

RE: buying BCP.

The best bet is to write directly to the current publisher, Walter de Gruyter, 200 Saw Mill River Road, Hawthorne, New York, 10532, USA. De Gruyter has only a few copies left and they probably don't send them to distributers. You might contact booksellers in England. My book was published there for some time, and was remaindered. I was asked if I wanted to buy 1700 copies for $1.75 each, and I was too poor to do it, oh woe is me. But some of those copies might still be in second-hand bookstores. American used-book sellers might also find copies.

Best to all,   Bill P.

Date: Thu Feb 27, 1992 11:19 am PST
Subject: Re: HCT details

[From Rick Marken (920227)]

Martin Taylor (920226) says:

> I feel that I am trying to provide substantial strengthening
>to the PCT model by showing how it follows from first principles, and the
>main protagonists react as if I were trying to dissolve their fragile edifice.
>Non comprendo.

I'm personally sorry if it seemed that way. I was just pointing out a fact about how a hierarchical control system works (the stuff about lower level reference signals being generally non-zero). I didn't feel defensive -- just trying to be informative. You say that this point is "unimportant in the general scheme of things" .. But "It could become important later". That's fine.

I did have a problem understanding the "degrees of freedom" point. You say:

>2. Neither Bill nor Rick accept the point that the degrees of freedom argument
>implies that there are uncontrolled but potentially controllable percepts.
> Point 2 is crucial, and we really should come to an agreement here.

I'm willing to try to come to an agreement here. I just don't understand your point -- really. I agree that there are an infinity of potentially controllable perceptions that are not controlled. I also agree that, given the limits on output degrees of freedom there is also a limit to the number of perceptions that can be controlled (this gets a bit closer to my understanding of the degrees of freedom problem in a hierarchy -- or any other arrangement--of control systems). I guess I just didn't understand what this has to do with control theory. I didn't see, for example, how this fact makes control theory a better model or the model of choice for living systems. I'm inclined toward control theory because it explains how organisms control. The theory is justified (and other theories ruled out) because it explains this phenomenon -- others don't.

>3. Bill brings up a new point that the boundaries of categorical decisions
>are unstable and context-dependent, and therefor that the continuous
>underlying percepts form the subject of control, not the decisions based on
>the categorical boundaries.

It seems to me that Bill has more than enough to do (he makes "working" look a lot easier than retirement) but he seems to manage to get it done -- so I'll let him deal with this if he wants.

> all the percepts in a stably
>controlling portion of a hierarchic control system will stay very close to
>zero. I lay awake many hours last night trying to see another stable solution
>to this organization, and failed. That doesn't mean there isn't one, but
>I can't see it at the moment. Bill and Rick must show that my condition is
>actually not a solution at all. Integration that is simply low-pass filtering
>does not do it. Non-linearities in the comparator responses might do it,
>but I have my doubts.

I think there must be some misunderstanding here. I don't seem to understand your problem. The best I can do (since I'm not a mathematician) is offer the spreadsheet as an existence theorm -- run it and you will see that you can achieve stability in a hierarchy of control systems with non-zero references. In fact, I think this is the ONLY way to achieve stability.

If ALL the reference signals going to a lower level specified zero levels for those perceptions, and those zero perceptions were brought about by the lower level systems, then all the higher level systems whose perceptions are a function of those now-zeroed perceptions would be zero (assuming that the higher level perceptions are linear functions of the lower level systems -- even if they were more complex functions they could only have one value because there is no way to change the components, which are always zero once stability is achieved).

Hasta Luego   Rick

Date: Thu Feb 27, 1992 5:46 pm PST
Subject: Zeros

[From Bill Powers (920227.1800)]

Martin Taylor (920226) --

> all the percepts in a stably
>controlling portion of a hierarchic control system will stay very close to
>zero. I lay awake many hours last night trying to see another stable
>solution to this organization, and failed. That doesn't mean there isn't
>one, but I can't see it at the moment. Bill and Rick must show that my
>condition is actually not a solution at all. Integration that is simply
>low-pass filtering does not do it. Non-linearities in the comparator
>responses might do it, but I have my doubts.

Try looking at this another way. Given a reference signal at a high level, what states of lower-level perceptions are required to satisfy it? This depends, of course on the perceptual function involved. Suppose we are talking about a perception of a relationship, like "circling around something" -- ice-skating in a circle around a fixed point. When the relationship is in the required state -- exactly in it -- the perception matches the "circling around" specification, but at lower levels a continuous series of actions is required to maintain it. From the skater's point of view, the scenery must be sweeping around at a constant non-zero velocity. The reference level for the rate of change of the visual image must be non-zero in order to maintain zero error in the sensed circling pattern. And to maintain the velocity error "at" zero, the legs must continuously be pumping to maintain the velocity and the curvature of the path.

I wish I had a better conception of how higher-level perceptual functions work. All I can say is that experientially, any perception seems to boil down to one signal that can get only larger and smaller. But the lower-level perceptions being combined to produce that signal may have to go through continuing complex dynamic, sequential, or logical space-time patterns to maintain that sense of one single thing continuing to happen.

As you noted rather a long time ago, converting an error in a system like this into an appropriate change in the pattern of lower-level reference signals must require a complex output function, something quite a lot more complex than a simple time-integrator. Yet when we look at behaviors involving such complex transformations, stipulating that the required functions exist and building their outcomes into our models as perceptual and output functions of simple kinds, we can still match an integral control system to the behavior. This says that control must involve gradual changes in the parameters of such complex pattern-generators, changes representable by a simple form of output function driving a variable pattern generator. For instance, moving the hand in a continous circle to track a target moving in a circle at varying speeds can be handled quite well by a control system with a leaky integrator in the output function. Such a model has to contain a circle-generator with only the velocity being variable, but if we just make a circle generator with a variable angular rate, that seems to do the trick.

The concept of a "zero" reference signal isn't necessarily easy to interpret. If there's a reference-relationship called "on", what condition of the two configurations corresponds to a zero reference signal? This isn't "not on;" it means that the sense of on-ness is to be avoided. Don't drop your cigarette "on" the floor, meaning "avoid creating the relationship cigarette on floor". This, I think, is generally the sort of meaning that zero reference signal or perception has at the higher levels -- zero means that the perception is not to exist. Any non-zero state of the perception implies presence of the condition being sensed, ranging from barely detectable or mostly incorrect to maximum (optimal?) presence.

Truly zero error is a rare condition. It's the amplified error that produces the behavior that is keeping the error as small as it is. The amplified error works by altering reference signals in lower systems. Even when the higher-level reference signal is zero, lower-level reference signals generally have to be non-zero in order to maintain the higher perception at zero.

One last point that might help: uncontrolled perceptions. As you've been pointing out, we do not control every perception. If a higher-level perception is composed entirely of controlled lower-level perceptions, then as you say the higher system will be protected against disturbances and will have little error-correcting to do. But if the higher-level perception is composed partly of controlled perceptions and partly of uncontrolled perceptions, disturbances can appear at the higher level even though the controlled component is perfectly controlled. A simple example is controlling a relationship between yourself and someone else, such as the distance between the two people. You have perfectly control over your own position, but no control over the other's position. As the other person moves, the relationship is disturbed; it is corrected by altering your own position.

Best   Bill P.

Date: Fri Feb 28, 1992 5:25 am PST
Subject: foot dragging

[From: Bruce Nevin (Fri 920128 07:22:15)]

(Bill Powers (920225.0800)) --

I liked this cockroach walkthrough a lot. But I found myself wondering how the critter starts out from a standstill--the old joke about the centipede being unable to decide which leg to move first. If it starts from a rearmost leg, then as that leg pushes the body forward do the other legs just drag along the floor?

So I took it out to behavioral outputs. Perhaps this connects with Greg's epiphany regarding HPCT vs non-HPCT perspectives (the model's point of view vs. the observer's, roughly?).

Quickly, as I'm more than ever under the gun today, here are freeze-frame shots of footpads on the floor, where + means the foot has just moved foreward (or is maybe the place a now lifted foot will land) and - means the foot is on the floor and the leg is in motion backward. The straight line is the starting point. Spaces count.

```
    |-  -  -                    |    -  +-
0   |                       4   |
    |-  -  -                    |    +-  +
    |                           |
    |                           |
    |  +-  -                    |      +-  +
1   |                       5   |
    |-  -  -                    |    -  +-
    |                           |
    |                           |
    |  -  +-                    |      -  +-
2   |                       6   |
    |  +-  -                    |      +-  +
    |                           |
    |                           |
    |    +-  +                  |        +-  +
3   |                       7   |
    |  -  +-                    |      -  +-
```

This crude approximation leaves out the "ripple" effect, as though the three +s at each stage (fewer in 1 and 2) were simultaneous. Nonetheless, it got me past the conceptual boggle about foot dragging. Also, it points up the observer's configuration perception of walking always with two legs on one side vs the middle one on the other, which is easy to see in any six-legged creature, such as a housefly. If Bill is right, this is a byproduct. But is it possible that something controls for synchrony of front and rear legs on the same side? Starting out would then look like this:

```
    |-  -  -
0   |
    |-  -  -
    |
    |
    |  +-    +
1   |
    |-    +-
    |
    |
    |  -  +  -
2   |
    |  +  -  +
    |
    |
    |    +-    +
3   |
    |  -    +-
```

A parsimonious construal of HPCT predicts (if Bill's analysis is correct) that we will find the first version, and not the second, in actual insects. What could motivate the extra level of control (for synchronization) if the simpler latch system works?

The external observer, catching on the the configuration (2 vs 1) of the established gait, would not attend especially to the first two stages to see whether it starts out with one hind foot, then the other hind and one mid foot, before getting into the swing of it. If they do notice, it probably gets interpreted as a kind of fidget or shifting of weight prior to actual walking. Simple observation could point up this crucial difference of perspectives. Quick! Find a cockroach!

And how do spiders do it, with 8 legs? Might be easier to observe there. And if they also evidence a latch-gait system, all the more convincing, given the genetic separation of insects from arachnids.

I'll get back, hopefully next week, on language, perception, and theories of grammar. No time now, and there are some deep problems here that are not easy to define.

Bruce

Date: Fri Feb 28, 1992 9:21 am PST
Subject: Re: Zeros

[Martin Taylor 920227 11:30] (Bill Powers 920227 10:00)

>One last point that might help: uncontrolled perceptions. As you've been
>pointing out, we do not control every perception. If a higher-level
>perception is composed entirely of controlled lower-level perceptions, then
>as you say the higher system will be protected against disturbances and
>will have little error-correcting to do. But if the higher-level perception
>is composed partly of controlled perceptions and partly of uncontrolled
>perceptions, disturbances can appear at the higher level even though the
>controlled component is perfectly controlled.

As so often happens, you pre-empt me. I had intended to bring this point up when I got the zeroing problem sorted out. I have been too busy these last couple of days to deal with it or with Points 2 and 3 of my earlier posting. I hope to be able to get back to the issue at the weekend.

I accept all of what you say about actions being required to maintain high-level stability. I find that my intuitions about how this works are self-contradictory.

Till later. Martin

Date: Fri Feb 28, 1992 9:47 am PST
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 neverthless 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 ot say about causal atributions 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.

Mary Powers:
I have two friends with whom I went to college who I believe would be very interested in reading the latest issue of Closed Loop as well as the one before (dealing with higher level system control). Should I send you $10 dollars with their addresses, or what? Do you still have copies of the previous issue? Let me know what would be best/easiest.

Carpe' Diem Mark Olson

Date: Fri Feb 28, 1992 2:56 pm PST
Subject: Auroral Activity Warning

[from Gary Cziko 920228.1600]

Unrelated to PCT, but because I care . . .

There is a good chance for auroral activity (northern lights) in North America the night of Feb. 29-March 1 (Sat. to Sun.). This solar flare was almost an order of magnitude greater than the one that was "supposed" to give us show last Wednesday night but fizzled.

In North America, potential viewing extends to 45 degrees north latitude (approximately north of a line from northern CA to NV to UT to CO to Northern OK to northern AK to TN to northern SC).

Hint: If you live south of the equator, look the OTHER way (do they call them Southern Lights down there?).--Gary

----
Gary A. Cziko Telephone: (217) 333-4382
Educational Psychology FAX: (217) 244-40538
University of Illinois Internet: g-cziko@uiuc.edu
1310 S. Sixth Street Radio: N9MJZ
210 Education Building
Champaign, Illinois 61820-6990
USA
