Date: Fri Oct 02, 1992 9:53 am PST
Subject: Progress, Diagrams

[From Dag Forssell (921001-1)]

Progress report:
Presented the "second day" program with leadership applications of PCT (Leadership, Vision/Mission statements, Teamwork, Performance reviews) yesterday. It was well received. Christine have strong signals that we will be invited back to teach other groups within the company.

Diagrams:

Martin: Thanks for direct post with comments on Red, Green and Blue cones from your correspondence with Bill. Do I recall that you mentioned the availability of your hypercard indexed files on PCT to Bill Silvert some months back? It occurs to me to ask you if you have submitted a set of your files to Greg Williams for historical purposes. He will keep anything confidential if you ask, but your correspondence with Bill will not be lost to the world if your house burns down.

I look forward to a report on your Paris experience.

Greg, thanks for quick and encouraging reply also. As discussed this am by phone, you posted direct so I would get it faster. I think your reply belongs on the net for anyone to see, so here it is with your permission:

Date: Mon Sep 28, 1992 10:03 am PST
Subject: Diagrams

From Greg Williams (920928 - direct)

Sorry I've been too busy to comment on your diagramming before now. I'll be gone tonight and tommorrow, at home Wednesday on.

>Dag Forssell (920926)

>I'll be happy to mail good looking charts of the final result to those
>deserving souls who help out.

And for the CSG archives, too, I trust!

>As I labor on this, I get uncertain about the levels where I have placed
>the different happenings of passing a car. Have I stretched this too
>much?

My overall suggestion, considering your motivation for doing the diagramming and the intended audience, is that you should not be worrying too much about whether every detail is perfectly worked out and whether there are any inconsistencies. I think your main point should be that this is the level of analysis necessary for a detailed description of the underlying phenomena involved in any behavioral scenario. Your managerial/engineering audience will get the message that the required analysis is anything but trivial, but that PCT is up (or is getting to be up) to the task. Remember, this whole exercise isn't to convince non-PCT psychologists that PCT is useful for understanding the genesis of behavioral events, but has the primarily heuristic function of relating the theory to "real life" events. At least that's the way I see it. If I'm approximately correct in my assumptions, the demonstration of how such an analysis would look is FAR more important than whether all the details are exactly correct. Besides, how would you verify (ooops, sorry -- falsify) the theoretical details at this stage of our knowledge about actual peoples' CNSs?

>Bill Powers (920926.1900)

>An ambitious undertaking. It makes me nervous, because it takes the 11
>levels as Gospel about 20 years prematurely.

Maybe 100 years or even more, with respect to falsification possibilities!

>Also, it doesn't (can't, really) capture the parallel nature of systems
>at these levels, and the branching networks that underly each system at
>each level.

Yet it gives the flavor of the depth of analysis required even WITHOUT getting into the parallel business.

>I've never seriously tried to fill in all 11 levels like this.

But see "A Cybernetic Model for Research in Human Development," in LCS [I].

>Dag Forssell (920927-1)

>My wording for the upper levels was hasty (and sloppy), pressing against
>time to post something for the sake of feedback. I had spent more time
>and felt more comfortable with the lower levels.

As I said above, I don't think that is a big problem.

>But this diagram can be supplemented by others (drawn to a smaller
>scale, (such as LCS I: page 278) designed to convey the parallel and
>branching nature of the model. I am working on that too.

Sounds good.

BP>>This sort of example isn't too difficult to construct at the lower
BP>>levels..

>I would like to do just this, then leave the upper reaches blank. My
>post yesterday will provide a start and a vehicle for communication on
>the net. The net result for now is simply to scale back my ambition.

How about the arm demo (new version almost ready) and/or the gatherings demo as explicit examples of lower levels operating in real time? There, you KNOW how the levels are organized!

>For purposes of Vision/Mission statements, which are a leadership
>application of PCT, I will be the first to suggest that people use the
>general concept of layering, but use labels that make sense to
>themselves as they see fit.

>By the time a person has created a statement for a business, using some
>of this terminology, it will likely appear to have more layers than we
>talk about in the model.

Quite frankly -- I've told Bill, of course -- I have never bought into Bill's hierarchy AT ALL except as a (not very well fleshed-out) existence proof for AN instantiation of PCT details. At this stage, I think PCT has a lot to learn from living organisms before pushing HPCT as THE instantiation. Your leaders are providing data which shouldn't be ignored by the theoreticians.

>Leaving the upper reaches of the model blank invites people to think
>this through and create layers of terminology that make sense to them.

Right on!   Best wishes,   Greg

>How about the arm demo (new version almost ready) and/or the gatherings
>demo as explicit examples of lower levels operating in real time?
>There, you KNOW how the levels are organized!

Right now my mind feels blank. Time to go back to bed and get some more sleep (after a very short night) before I start in earnest to complete the workbook with diagrams for the third day. In a few days, I will study these programs so I can show them off next Wednesday. In the meantime, can you make any suggestions on the arm demo and crowd using my format?

Thanks for help!

Dag

Date: Fri Oct 02, 1992 10:42 am PST
Subject: Re: Instructional Labs

>[From: Clark McPhail (921001)]

>>>Dennis Delprato (920930)--

>I do feel that for purposes of instruction, it is a good policy to
>keep things as simple as possible so as not to obscure the really
>big points of PCT with matters that are best considered by the
>more advanced student. Bill's Demo 1 and Demo 2 (not currently
>in a form suitable for labs) nicely illustrate a gradual
>build up in complexity.

Dennis:

I'm not sure to what laboratory purpose(s) you want to put Bill's Demo 1 and Demo2. I have put Demo 1 on a file server for up to 50 DOS machines. Students work through the entire set of exercises before returning to one of their choice which they then run a second time. Having done so they print the screen display of the traces for that run and take that print away with them. They then answer a series of questions on the printout and submit their answers with the printout for my evaluation.

Clark

Date: Fri Oct 02, 1992 10:43 am PST
Subject: Re: Arm (to B.P.); Influence Summary

FR Clark McPhail (921001_

>>RE Greg Williams (920928)

>An aside: What's going on in HYPNOSIS in PCT terms???

I am not surprised that someone concerned with "influence" would be concerned with the "ultimate" form of "social control": hypnosis. My impression for years was that hypnosis works because the subject focuses upon and tells him/herself to do exclusively what he/she is instructed to do by the hypnotist. Then I discovered with work of Theodore X. Barber, a psychologist who is now retired but who examined the phenomenon of hypnosis across a forty year span of time. Simply stated, Barber rejects the "trance state" theory of hypnosis and has generated a considerable body of empirical evidence supporting his critique. Further, Barber has advanced an alternative interpretation which turns the commonsense view of hypnosis on its head. Instead of the subject-as-passive receptacle of the hypnotist's suggestions, Barber construes subjects as exercising variable degrees of imagination; that is, they are variably capable of imagining the outcome the hypnotist "suggests" and then carrying out the actions require to fulfill what they have imagined. In one experiment he compared the performances of a standard repertoire of "hypnotic feats" (e.g., the plank posture, arm levitation, locked clasped hands, verbal inhibition, posthypnotic response, selective amnesia and analgesia, etc) for 40 subjects who had been through a standard 15 minute trance induction procedure, 40 subjects who had been through a brief sequence of "positive suggestions" ("You can do these things and it will be a lot of fun. . ."), and 40 "control" subjects. The subjects in the "positive suggestions" condition exceeded the performances of the subjects in the "trance induction" condition (and the control subjects as well). A synopsis of many of Barbers experimental papers, which numbered over 150 at last count, of his critique of the hypnotic trance state hypothesis, and of his alternative interpretation of hypnosis, are found in his (1972) "Suggested ('Hypnotic") Behavior: The Trance Paradigm vs. An Alternative Paradigm." Pp. 115-182 in Erica Fromm and Ronald E. Short (eds.) _Hypnosis: Research Developments and Perspectives_, Chicago: Aldine. (You will find a short synopsis of Barber's work and a limited set of references in my (1991) book, _The Myth of the Madding Crowd_. This is not a "PCT" interpretation of hypnosis but it leans in a direction that might be recast in such terms.

A related aside. What is going on in the hierarchy of perceptual control systems when individuals ingest small, moderate, large amounts of alcohol? Heroin? cocaine? marijuana? Do these simply affect the operation, e.g., the loop gain, of every control system or might these turn off higher level systems? I would appreciate any comment on these matters (and on hypnosis) from anyone on the loop.

Date: Fri Oct 02, 1992 11:10 am PST
Subject: Education as Influencing Reorganization

[from Gary Cziko 921002.0200 GMT]

Relevant to the recent discussion of purposeful influence and reorganization is an excerpt (sans emphases and footnotes) from a chapter of my book in preparation. The book is tentatively entitled _Without Miracles: Universal Selection Theory and the Second Darwinian Revolution_ and the chapter is "Education: Transmission of Truth or Growth of Fallible Knowledge?"

I currently feel that it is indeed possible for people (e.g., teachers) to influence long-lasting changes in others' (e.g., students') control systems. Although the following was not written to make this point, I think it is compatible with it. Reactions (especially from Greg Williams and Bill Powers and Chuck Tucker) are eagerly anticipated, but I don't see myself engaged in a protracted Powers-Williams type discussion.

==

To apply the perceptual-control-theory notion of reorganization to education, let us use an example--a person learning to swim. In its most rudimentary form, being able to swim can be defined as staying alive in water which is deeper than one is tall, that is, being able to "tread water." One way to "teach" a non-swimmer to swim is the throw the person into a body of deep water (we could call this the "immersion" method swimming instruction). This will likely create error in the student since she will have difficulty keeping her head above the water in order to breathe. This perceived error in a crucial variable will trigger reorganization so that the student will immediately begin to vigorously move her arms and legs in random patterns in order to find some way to maintain her nose and mouth above the water's surface. If a behavioral pattern (actually a perceptual-behavioral control loop) is found which allows her to breathe (if even only a few gasps before disappearing below the surface again), the randomness of the movements will decline until the student is able to constantly keep her head above the water, at which point we would say that she has learned to swim. In effect, the student has now gained control over a variable which she could not control previously and so by our definition learning has taken place.

Since the student did not initially know how to swim, the initial movements were of necessity blind attempts to swim. But while the student did not know how to keep her mouth above the water, she could perceive how successful she was in her attempts (getting her eyes above the water is better than below, but not quite good enough). This then provided a criterion for selection among the various behavioral patterns attempted and allowed the student to learn from her mistakes, eliminating those patterns of movement which did not succeed in getting the head above the surface and retaining those which did.

We can also easily imagine that the learner in this example would be very highly motivated since failure to learn could result in death. From a perceptual-control-theory perspective, motivation simply refers to error (that is, a difference between a perception and the reference level for that perception) which results in action to eliminate the error (see Figure 7.1). From this perspective, motivation is considered to be internal to the student since the reference level of the controlled variable is determined by the student, not by the environment.

We need to point out, however, that this immersion method of swimming instruction may well fail for any particular student since there is no guarantee that the student will come up with an effective control system for treading water within the few minutes available before lack of oxygen leads to loss of consciousness and death. Clearly, a less drastic approach to swimming instruction is needed. There are a number of ways in which this method could be improved. First, we could simply allow more time for learning to take place. This could be accomplished by allowing the student to practice at the edge of a swimming pool so that she could reach out and hold onto the edge of the pool when she felt herself going under water. Or she could be allowed to practice in water that was only neck deep so that she could simply stand in the water at any time to breathe. Given more time to try out new patterns of movement and eliminate those which are ineffective, the likelihood of successful learning would be increased.

Another approach to facilitate learning would be to attempt to accelerate the learning process itself using verbal instructions ("move your hands horizontally in the water from your sides to the front and back again") or demonstrating a model for imitation, or a combination of the two ("do it like this"). Such instruction could be useful in constraining the attempts made by the student (for example, the student would not now attempt vertical movements of the hands). But such instruction, no matter how effective, can in no way transmit the skill to be learned from teacher to student. Even if the teacher provides a model to imitate, the student must still learn on her own how to imitate the model. The perceptions the student has of the teacher demonstrating a technique are very different from the perceptions the student will have when she is able to successfully perform the technique herself (watching someone else swim is a very different experience than that of actually swimming oneself). Models and instruction can provide useful information in the form of constraints of what not to try, but they cannot provide explicit instructions concerning exactly what to do.

In addition to allowing more time for the learning to take place and providing constraints in the form of models and verbal instruction, the teacher can also provide easier access to the knowledge or skill by providing a series of less demanding intermediate goals. One way is to break down the skill into a number of subskills and provide opportunities for the subskills be acquired. In the swimming example, the teacher could have the student stand in shoulder-depth water and have her make horizontal movements with her arms until she feels an upward force lifting her weight from her legs. After this is mastered, the student could learn to move her legs by holding onto a float and kicking her legs until she feels herself rising from the water. After practicing the arm and leg movements separately, she could then attempt to combine them, first in water shoulder high and then in deeper water.

By breaking down a complex problem into easier subproblems learning is facilitated since the probability of finding a solution to each subproblem is higher than that of finding a solution to the more complex problem--success in learning to make effective arm alone in swimming is more likely than success in learning to make both arm and leg movements together. A selectionist-reorganization view of learning sees the teacher as constantly aware of the student's current abilities and continually imposing upon her problems which are just a bit beyond these abilities. Assuming that the student wants to be able to gain control over this new situation, reorganization will take place until control is achieved at which time new demands are imposed (after learning to tread water, the breaststroke is attempted; after addition is learned, subtraction is introduced). Such a view of learning is consistent with Russian psychologist Lev Vygotsky's (1896-1934) notion of the "zone of proximal development" in which the student is able to try and eventually successfully master new problems which are beyond her independent capabilities but which can be learned with the assistance of a teacher. Note that the teacher's role in this learning is not one of a transmitter of information or knowledge but rather one who provides support to the student and arranges the student's learning environment in such a way that the student is continuously challenged by situations and problems which are just a bit beyond her current competence. In other words, the teacher arranges the environment so that the student is continually encountering error, but error that is not too large so that the student's reorganizing efforts are likely to be successful and set the stage for the next introduction of error. This view is also consistent with the idea now popular in education that a successful teacher provides educational "scaffolds" for his students. Such scaffolds can be seen as teacher-provided support platforms which provide support in breaking down complex physical and cognitive problems into more easily mastered subproblems. While we used the physical skill of learning to swim in this example, all that has been said is also applicable to other more cognitive skills such as those involved in learning mathematics and developing reading comprehension and writing skills.
===
P.S. The idea of using swimming as an example was inspired by the Robertson and Power's text. Gary A. Cziko

Date: Fri Oct 02, 1992 11:53 am PST
Subject: Misc subjects

[From Bill Powers (921002.0600)]

Postcard:
Back from a trip with Mary that stitched back and forth and up and down New Mexico for four days. Saw the Very Large Array again, and one radio telescope of the Very Long Baseline Array in Pietown, NM. Saw them both move this time. Visited the observatory I designed in Las Cruces -- locked up with nobody around, as it's run by amateur astronomers now, but it's still a beautiful setting. No signs of Georgia O'Keefe at Abiquiu, but you can see her colors everywhere in the Chama River valley. Camping at the end of September gets a little chilly at night. There's a lot of geology in New Mexico and you can usually see about 40 miles in any direction, on the ground, because the valleys all rise at the edges. Interesting place where we humans live.

Dag Forssell (920927-1) --

Added suggestions for System Concept level: entities, organizations, persons, realities, disciplines.

Cliff Joslyn (920922.1900) --
RE: Turing Test and intelligence.

>The TT helps us understand when we're PERCEIVING intelligence. But
>like redness, that perception is intuitive, and automatic; and
>frequently faulty.

It's also parochial -- that is, academics tend to rate verbal skills high, while others consider a preoccupation with words (as opposed to, say, financial manipulations) rather stupid. I don't think the TT helps us to understand anything but the problem, which is that a word like "intelligence" can't possibly wrap up all the dimensions of human Being.

Greg Williams (920928) --

I agree with

>1. A disturbs particular perceptions being controlled by B so that
>B compensates for the disturbances with actions which A wants to
>perceive.

>2. A arranges B's environment so that when B controls for
>particular perceptions, A perceives what he/she wants to perceive.

but I have a problem with

>3. A arranges B's environment so as to trigger learning
>/reorganization in B's control system resulting in actions which A
>wants to perceive.

>4. A applies physical constraints or threatens to apply physical
>constraints to B so that B's actions are as A wants to perceive.

According to my model of what triggers reorganization, these would both mean arranging the environment so that B suffers critical error (you notice my return to Ashby's term) such as hunger, thirst, pain, illness, suffocation, "stimulus deprivation," or whatever you want to put on the list. By definition, reorganization is unsystematic. This means that you can't predict what behavior will be used to correct the error unless you have removed all means of correcting it but one, which is within B's capacity to learn. That's easy to do with a lower animal or a child, but hard to do with an adult human being. You do note that these methods involve conflict, but you don't mention that the outcome is largely unpredictable because reorganization is involved.

In the second section, I don't understand

>3. Onset of learning/reorganization at a particular time is a
>function of reference signals, input/output functions, and
>environmental disturbances at that time.

>4. The path of learning/reorganization is a function of (possibly
>randomly generated) successive sets of changes in reference signals
>and/or input/output functions ...

It seems that you're allowing for systematic reorganization here. Why isn't that just the operation of a higher level control system, which itself has to be learned?

>5. Whether or not the criteria for ceasing learning/reorganization
>are met by the reference signals and/or input/output functions at
>any point on the path of learning/reorganization is a function of
>reference signals, input/output functions, and environmental
>disturbances at that point.

Are you proposing some particular mechanism here, or is this just in general? What is it that judges whether the criteria are met? How does that judgment affect the continuation or cessation of reorganization? I can't visualize the arrangement you're talking about. Can you boil it down to a specific model?

>6. At any time, the criteria for ceasing learning/reorganization
>are functions of reference signals and input/output functions at
>that time.

The CRITERIA are functions of reference signal and input/output functions? Now I'm thoroughly confused. How do these things affect the criteria? Do you have a mechanism in mind?

The main motive for my simple model of reorganization was a need to explain how animals learn such things as walking in a figure eight to get food -- situations where what is learned has no necessary or a priori connection with the reasons that it has to be learned. I then realized that NOTHING has a neccessary or a priori connection with the need for learning a specific behavior. Even learning to eat certain items having particular appearances or smells or tastes has nothing guaranteed to do with assuaging hunger or correcting the underlying nutritional state. So that's where random reorganization based on critical errors came from.

It seems to me that you're proposing something different here. What is it? What phenomena of learning does it explain that my version of reorganization doesn't explain? And how does your explanation work?

Martin Taylor (920927.1830) --

I said:
>>The "impedance" concept is sort of ingenious, but I can't see how
>>to model it so that a particular output from a control system
>>would be spread out among all the different possibilities -- what
>>would keep all of them from trying to happen at once?

You said:
>Well, I must have misinterpreted you some months or more ago,
>because I thought it was your idea. It's "always" been a part of
>my concept of PCT. Anyway, if it wasn't your idea, here's how I
>see it. They all DO try to happen at once, but they can't. They
>inhibit one another.

I did speak of impedance-matching in a different context, having to do with power gain. But the idea of "competing behaviors" inhibiting each other doesn't seem very plausible to me. How does one behaving subsystem know which other behaviors to inhibit and which to leave alone? I don't think you could make a runnable model out of this, but you're welcome to prove me wrong. Your idea isn't even self- consistent, because you follow the above by saying

>It is the world that stops them all happening at once. If they
>could, they would.

But that means that they are NOT mutually inhibiting each other inside the system. They're all in a state of perpetual conflict. This seems like a lousy design for a control-system hierarchy where you would like the loop gains all to be quite high.

>Outputs are going in all directions, but the world prohibits some percepts
>from actually being controlled. So the "taking bicycle" percept/reference
>cannot be satisfied if you are sitting in the car.

So when you're driving along in the car, you're wishing you could be riding your bicycle, and vice versa? And how could your outputs be going in all directions while you're using them to drive a car?

>I guess in the background there is another point I meant to discuss
>at some time--"giving up." When there is persistent error, one
>possible and often used response of an ECS is simply to reduce its
>gain to zero, to give up on a hopeless situation.

I have already brought up this idea on the net, but not by supposing that every ECS can perceive and judge "situations" as being "hopeless." My solution was a comparator that has a curve that I drew like this:

```
     ^                  *      |
     |                *    *   |
Error|signal      *         *  |
            *                * |
*                   -         *|             +
-------------------------------*-----------------------------------
         <-- actual error-->   |*                           *
                               | *                *
                               |  *          *
                               |   *     *
                               |      *
```

In the central region, feedback is negative. When a disturbance gets so large that the peak of the curve is reached (in either direction), the error signal begins to fall for further increases in disturbance. This results in a drop of output, still more error, still less error signal, and so on to the "giving up" regions at the ends. I had previously proposed this as a possible explanation of operant conditioning under conditions of high deprivation. The giving up process is reversible; if the disturbance falls, the error will drop and the error signal will rise, producing more output, until the system passes over the peak of the curve and snaps into the region of negative feedback control again (the whole system may be unstable in a region near the peaks, the size of the region depending on the loop gain).

>When two ECSs are in conflict, it is impossible for both to bring
>their errors to zero simultaneously. That's almost a definition of
>conflict.

That's the definition I've always recommended and used.

>If one ECS, for example, has a reference to perceive the body as
>bicycling to work, and the other has a reference to perceive it as
>driving, there is no compromise intermediate position (possibly a
>motorcycle, but let's suppose none is at hand).   The mechanism is
>that one ECS gives up attampts to control. Either the person
>cycles, or the person drives, but not part of each.

Unless you use something equivalent to my comparator function above, you will require that an ECS do something other than control its own perception. It seems simpler to me to invoke a higher-level system that selects one of several means of locomotion (the kind of explanation you used for explaining why the wastebasket wasn't moved on the third day). This has the advantage of not pitting the outputs of high-gain (competent) control systems against each other, which always leads to a lessening of the capacity to control if not its complete destruction. I don't think you're taking full advantage of the concept of hierarchical control, at least not consistently.

>My claim is that The Test is always ambiguous. P can tell that Q is
>controlling for a percept that incorporates some CEV that P has
>disturbed, but P can never tell that the percept Q is controlling
>for corresponds exactly to the CEV that P's percept corresponds to.

One application of one hypothesis leaves the Test ambiguous. There is nothing to prevent you, however, from trying out more hypotheses aimed at reducing the ambiguity (as you pointed out to your colleague). I don't think there is any more ambiguity in applying the Test than there is in guessing at the causes of any natural phenomenon. A single measurement is always ambiguous. One employs strategies aimed at eliminating alternative hypotheses until no more alternatives can be found. Then you go with what remains. Your statement that you then resort to statistics is gratuitous; you may or you may not, depending on how unclear the hypothesis is when you're done and whether you happen to like statistics. Lots of people will think up one explanation on the basis of the flimsiest evidence and assume it's absolutely correct. They wouldn't do any better using the Test.

RE: applying disturbances (stuck phonograph record division):

>Yes, ideally any observer should avoid disturbing the thing
>observed. But any disturbance to a controlled variable causes
>error in the controller, even if momentarily. The tester is
>inevitably controlling for perceiving a change in a variable, or a
>resistance to change (if the circumstances are appropriate).

The disturber should NOT control for a visible change in the variable being disturbed. That simply creates conflict. What the disturber should do (where possible) is alter some OTHER variable that is loosely coupled to the putative controlled variable. This will elicit an opposing change in the controller's actions even if the controlled variable doesn't visibly change. In low-gain situations (like the coin game) this doesn't matter so much. But when the control system involved is a very good one, insistence on seeing the controlled variable actually change will result in applying very large forces to the controlled variable, with a probable change in what variables the controller is controlling. The Test is most accurate when the controlled variable doesn't change at all (that you can see).

Dennis Delprato (920929) --

Request for lab experiments noted. In the queue.

Bill P.

Date: Fri Oct 02, 1992 12:18 pm PST
Subject: Re: Why 99%?

From Tom Bourbon [921002 -- 10:50 CDT]

Subject: Why 99%?

>Martin Taylor 920929 16:00

>My presumption is that you get the 99% prediction because the subsystems
>(perhaps ECSs) that are involved in the task are those that support
>very many different kinds of behavior, and so are not readily disturbed
>by contextual differences.

What do you mean by "contextual differences?" I believe your presumption is wrong, but I am not certain what you are saying. By "contextual differences" do you mean that target positions follow different random paths on every trial and the cursor is disturbed by a different random function on every trial? That is true, but how does that fact in itself explain the success of the model? The "subsystems" that "support" tracking certainly "support" many other controlled results of action, but that fact in itself does not explain .99+ correlations between predictions and results. The same 1st-level control loops "support" all of our actions and all of our controlled perceptions. How? Why? The answer cannot be merely because they do.

****

In reply to those remarks by Martin Taylor 920929 16:00, Greg Williams wrote: Greg Williams (920930 - 2)

>I believe that the highly precise predictions which have been achieved
>by some PCT tracking models are due to the condition of "good" control

>being satisfied in the experiments to which the predictions are applied.
>That is, the subject is able to track close to the target (because
>the disturbances are "easy" to compensate for).

It was good to see that you wrote ""good" control" and ""easy" to compensate for." The quotation marks reveal your awareness that accurate predictions by PCT do not depend on tracking tasks that are so easy people can do them in their sleep. Since many people on the net have not read the tracking literature, or performed any of the tracking tasks, they might not share your insight.

The PCT model does not succeed by predicting that people will do extremely well on the task. Sometimes they don't. (I will give some examples later.) The way most of us (most of the three of us who do modeling) use the model is to let a person, or more than one person, do the task. Then we estimate at least two parameters (reference signal and integration factor) that produce a least squares fit of the model's simulated positions of the cursor and the person's positions during the first run. If the person did well, the model does well in simulation; but if the person did relatively poorly, so does the model. That is for high correlations between the results of model tracking and person tracking. It is not that the model "does good" and fits only if people "do good." Rather, the model does as well or poorly ad the person.

Then we often let the model track again, with new target conditions and never-before-experienced disturbances acting on the cursor. Then the person (or more than one person) does the task. If the person did relatively poorly on the first run, chances are that will happen on the second, disturbed, run. And the model will have the same problems. The agreements of results for model and person are often very detailed - - particular handle movements at certain points in the run that lead to very specific deviations of the cursor from the target. Let me cite a few examples from some of our work,

In the paper on 104 replications of tracking and predictions that I published along with several students (in Perceptual and Motor Skills, 1990), the range of correlations between 1800 pairs of positions of cursor and target (produced by the people during the second, predicted, runs) was .822 to .998. No one did terribly, but some were obviously less effective than others. The differences were easy to see during the runs and in the printouts of results. At the same time, correlations between model-predicted positions of the control handle and the persons' positions ranged from .989 to .998. Below are a few comparisons of cursor-target correlations and the associated correlations between model-person handle positions:

cursor-target	model-person
.834	.997
.920	.990
.822	.993
.998	.997
.974	.989
.987	.998
.996	.991

This list could go on for 104 replications, but I think you can see that the "best" and "worst" fits between model and person do not always correspond to "best and worst performance" by the person. A .993 agreement between model and person came on the "worst" of 104 replications.

I will grant that everyone did "pretty well" on the task, but then PCT is intended to explain and predict control, not something else. Faulting PCT for doing what it is intended to do would be like faulting Ohm's Law for describing relationships among voltage, current and resistance, but not describing last month's Dow Jones Averages on the New York Stock Exchange.

In my publications on two-handed tracking (whether one person using two hands or two people interacting) I report results very much like those summarized above. With sufficiently rapid changes in target positions and disturbances, and with some of the interactions that I create between the two control devices, I defy anyone to say the tasks are "easy." People often demonstrate, then report, feelings of extreme effort and tension. Good prediction does not require near-perfect performance or easy tasks.

That point is made even more clearly in data gathered by one of my thesis students, Wade Harman, in 1991. Wade set up a tracking task with four-degrees of freedom. An arrowhead-shaped target moved on the screen in X and Y, rotated in A and changed size in Z. Participants used a mouse to move the similarly-shaped cursor in X and Y, and a joystick to rotate it in A and change its size in Z. A few people found the task relatively easy; a few others thought it a high-tech form of torture. Control in X and Y was easy for nearly everyone, but Z was sometimes a challenge and A was agony for some people. Below are a few examples of correlations between cursor and target, on each degree of freedom) and the corresponding correlations between model and person for that variable:

X		Y	
c-t	mod-pers	c-t	mod-pers
.997	.976	.998	.998
.987	.992	.988	.989

Z		A	
c-t	mod-pers	c-t	mod-pers
.976	.939	.987	.984
.716	.828	.987	.984
.759	.870	.463	.880
.935	.956	.404	.954
		.430	.930
		.986	.989

Now we see some model-person correlations below .99, but they still exceed by orders of magnitude the modal and median correlations in the behavioral literature. Performance on Z could be very shaky, and on A it could be horrible, but the PCT model, which assumes control, was not shot out of the water.

I agree that the lowered level of agreement between model and person suggests a "problem" with the model -- we used four independent models, one for each degree of freedom in the cursor and target. We take the results as a suggestion that some degree of "coupling" or interaction between the models might be necessary. From our observations of participants, and our own introspective experiences while performing the task, we believe there is some "task switching" in which people learning the overall task take some time out from controlling A, allowing it to drift through a wider range of discrepancy before they act. But is that task switching? Couldn't we model it as an independent control loop for A with a lower gain or error sensitivity? Those are questions we will sort out next. I believe our work on this problem is in the spirit of Greg's remarks when he said:

>...the important question (if you already believe in PCT) is how
>to choose BETWEEN different negative-feedback models, and the only
>way to do this WITH HIGH PRECISION is by looking at transient (temporarily
>"poor") control. The "good" control case data simply don't allow
>sufficient sensitivity for comparing different models.

I like your emphases that indicate the relativity of "good and bad" control and of "high and low" performance. Your reasoned treatment is better informed than that of people who merely assert their beliefs and prejudices on these matters. PCT models control. When control is present, PCT predicts with precision -- even if the control is mediocre. It predicts with great precision if the control is precise.

Best to all, Tom Bourbon

MEG Laboratory	PAPANICOLAOU@UTMBEACH.BITNET
1528 Postoffice Street	PAPANICOLAOU@BEACH.UTMB.EDU
Galveston, TX 77550	PHONE (409) 763-6325
USA	FAX (409) 762-9961

Date: Fri Oct 02, 1992 12:27 pm PST
Subject: pct, medicine; questions

[From Francisco Arocha, 921002; 1202)

Question: Is there a reading (a study or a theoretical presentation) of higher levels of perceptual control? The problem I'm having is that as I try to understand higher levels in the hierarchy, my understanding becomes more confused and I have trouble thinking about an experiment in which I could achieve the degree of precision that is found in lower levels (the 99%). Do I have to settle for (much) less? The study that I'm planning (on which some ideas below) I think would be about category level and above. But I found these less precisely defined.

A second question: Is there any in the net familiar with the situated literature, aside from the situated AIers like Chapman & friends? I am referring to those in that camp who work on "higher" cognition in psychology and education (Greeno, Seely Brown, William Clancey).

******

PCT and medicine II

As I posted several weeks ago, I'm planing a PCT study in medicine. I have looked at some verbal protocols with the hope that they help me understand what the physicians think about. Also I have read some medical problem solving textbooks and atended a few clases, to see what is that medical students are taught, and medical ground rounds (where doctors discuss their cases). This I have done to come up with some hypotheses. Among the things that I have noted that seem interesting are the following:

1) there is a lot of disagreement among physicians about most cases. This disagreement covers the diagnosis itsef as well as the proper treatment (even if they agree on the diagnosis).

2) the students are taught some rules of thumb that seem to play an important role in diagnosis. These are things like "common things are common", "when you hear hoofbeats think of horses not zebras".

3) when doctors or med students give their explanations about a case, some are more interested in getting the diagnosis right whereas others are more interested in not missing any fatal disease, to the expense of being right (actaully, more accurately would be to say that some physicians mention that what drives their thinking is not to miss an important disease, whereas others say nothing about this and focus on producing the correct diagnosis).

4) Some physicians are very sparse in their explanations while others are very detailed. As far as I can see, there is no connection between this and getting the right diagnosis.

5) When physicians justify their thinking they focus on different signs and symptoms. By this I mean that one may mention that what "triggers" their diagnosis was, say symptom A, whereas another says symptom B. Again, there does not seem to be any connection between accuracy in diagnosing (or choosing the "right" treatment) and what they focus on.

These observations are informal; just by talking with physicians and med students. I still think that, at least for now, is to do the study in artificial situation rather than with real cases and real patients. The reason for this twofold: that I would not know how to do it and that I think doing it the other way would take me longer than the time I have (my postdoc ends in june).

I thank David Goldstein and Bill Powers for the comments to my earlier post. I hope that you both (and any other PCTer) can help me along the way. I would really like to complete the study.

Saludos,   Francisco

Date: Fri Oct 02, 1992 12:40 pm PST
From: g cziko
          MBX: g-cziko@uiuc.edu
TO: * Dag Forssell / MCI ID: 474-2580
Subject: Intro Document

Dag: Attached is the document I currently send out to new CSGnet subscribers. Let me know what you think should be added.--Gary

==
INTRODUCTION TO THE CONTROL SYSTEMS GROUP NETWORK (CSGnet)

This introduction to the Control Systems Network (CSGnet) provides (Control Systems Group net) provides
information about:

Why you might want to read CSGnet
Our subject matter: The control paradigm
The purpose of CSGnet
CSGnet participants
The evolution of the control paradigm
How to obtain text and program files
How to ask effective questions
Demonstrating the Phenomenon of Perceptual Control
The Control Systems Group
Literature references

WHY YOU MIGHT READ THE CSGnet

If you are curious about things that are new and exciting...
If you are dissatisfied with the explanations (or the lack thereof) in many of the "soft" life sciences and would like a more rigorous approach that has more power of explanation...
If you insist on thinking things through for yourself rather than accept what the establishment feeds you....

+++
OUR SUBJECT MATTER: THE CONTROL PARADIGM

Human control is the primary subject of CSGnet, but all forms of control are game. Here is a brief introduction by the primary creator and promoter of the application of the control paradigm to living systems, William T. Powers:

There have been two paradigms in the behavioral sciences since the 1600's. One was the idea that events impinging on organisms make them behave as they do. The other, which was developed in the 1930s, is PERCEPTUAL CONTROL THEORY (PCT).

Perceptual Control Theory explains how organisms control what happens to them. This means all organisms from the amoeba to humankind.

It explains why one organism can't control another without physical violence.

It explains why people deprived of any major part of their ability to control soon become dysfunctional, lose interest in life, pine away and die.

It explains what a goal is, how goals relate to action, how action affects perceptions and how perceptions define the reality in which we live and move and have our being.

Perceptual Control Theory is the first scientific theory that can handle all these phenomena within a single, TESTABLE concept of how living systems work.

William T. Powers, November 3, 1991
++

THE PURPOSE OF CSGnet:

CSGnet provides a forum for development of PCT in considerable detail, applications and testing of PCT and the dissemination of PCT to any and all who have a sincere interest in how organisms work.

CSGnet PARTICIPANTS

Many interests and backgrounds are represented here. Psychology, Sociology, Linguistics, Artificial Intelligence, Robotics, Social Work, Social Control, Modeling and Testing. All are represented and discussed. A challenging quality of participants on this net is that most are prepared to question and re-consider what they think they know, even if it requires that a LOT of previous learning be rejected.

THE EVOLUTION OF THE CONTROL PARADIGM

The PCT paradigm originates in 1927, when Harold Black invented the negative feedback amplifier, which is a control device. This invention led to the development of purposeful machines. Purposeful machines have built-in intent to achieve consistent ends by variable means under changing conditions. Examples are the heating system in your home, which keeps the indoor temperature constant despite the changing seasons and opening doors and the cruise control in your car, which keeps the speed constant despite changing road conditions.

The first use of this concept to better understand people was suggested in 1957 in a paper entitled "A General Feedback Theory of Human Behavior" by McFarland, Powers and Clark. In 1973 William T. (Bill) Powers published a seminal book called "Behavior: the Control of Perception," which still is the major reference for PCT. See literature below.

This book spells out a complete model of how the human brain and nervous system works like a living perceptual control system. Our brain can be viewed as a system that controls its own perceptions. This view suggests explanations for many previously mysterious aspects of how people interact with their world.

Since 1973 an acceptance of Perceptual Control Theory has begun to emerge among a few psychologists, scientists and other interested people. The result is that an association has been formed (the Control System Group), several books published, this net set up for communication and that a dozen professors are teaching PCT in American universities today.

HOW TO OBTAIN TEXT AND PROGRAM FILES

A number of documents and computer programs are available on a fileserver maintained by Bill Silvert. Although it is possible to obtain these files via e-mail, it is far easier to obtain binary program files via anonymous FTP. The Internet address for the machine is BIOME.BIO.NS.CA. CSGnet files are kept in the subdirectory pub/csg. Here is a listing and brief description of some of the files available.

```
dem1a.exe  Powers's demonstration of the phenomenon of control;
                  self-extracting archive for MS-DOS + mouse
dem2a.exe  Powers's demonstration of the control theory model
                  self-extracting archive for MS-DOS + mouse
biblio.pct  Williams PCT Bibliography; Text
blindmen.doc  Marken Paper 1992; Text
marken.bhx   Marken spreadsheet of hierarchical control;
                    Lotus spreadsheet in BinHex form for MS-DOS
marken.doc   Marken paper describing spreadsheet; Text
marken.wk1  Marken Spreadsheet Model; Lotus format for MS-DOS
```

NOTE: Any file not indicated as text should be transferred as a binary file.

Any TEXT file can also be obtained via e-mail by sending a request to the address SERVER@BIOME.BIO.NS.CA. For example, to get Williams's bibliography, send a message to the SERVER and include this command as the first line of the message:

get biblio.pct

The file will then be sent to you via e-mail.

HOW TO ASK EFFECTIVE QUESTIONS

Since PCT puts much conventional, well established wisdom on its head, it is very helpful to begin by demonstrating the phenomenon of control to yourself and studying a few references. It is helpful to study systems and control in general in addition to the texts that focus on PCT. As you catch on to what this is about, read this net and follow a thread that interests you for a month or more.

When you ask a question, please consider that in order to give you a good answer, a respondent will need to put your question in context.

Therefore, please introduce yourself with a statement of your professional interests and background. It will be helpful if you spell out what parts of the demonstrations, introductory papers and references you have taken the time to digest and what you learned.

People on this net are in various stages of learning and understanding PCT. When you get a reply to your post, please consider that the respondent who found your question of interest and invested time in a reply, may benefit from knowing how you perceived the answer. Did it answer your question? Was it clear? Were you able to understand it?

DEMONSTRATING THE PHENOMENON OF CONTROL

The phenomenon of control is largely unrecognized in science today. It is not well understood in important aspects even by many control engineers. Yet the phenomenon of control, when it is recognized and understood, provides a powerful enhancement to scientific perspectives.

It is essential to recognize this phenomenon before ANY of the discourse on CSGnet will make any sense.

Please download the introductory demonstration dem1a.exe.

THE CONTROL SYSTEMS GROUP

Serious enthusiasts of PCT have formed the Control Systems Group. This group meets once a year (1992: July 29-Aug 1) in Durango, Colorado, for informal presentations and exchanges. The group also publishes threads from this net. For membership information download the file csg.doc (not yet available as of June 11, 1992; soon to be).

LITERATURE REFERENCES

For a complete list of CSG-related publications, get the file biblio.pct from the fileserver as described above. Here are some selected, books on perceptual control theory.

Powers, William T., Behavior: The Control of Perception. Hawthorne, NY: Aldine DeGruyter, 1973, 296 pages. The foundation of PCT! A seminal book.

Robertson, Richard J. and Powers, William T., editors. INTRODUCTION TO MODERN PSYCHOLOGY; The Control Theory view. Gravel Switch, KY: The Control Systems Group, 1990, 238 pages. Textbook on psychology for universities. Highly recommended.

William T. Powers, LIVING CONTROL SYSTEMS: Selected Papers. Gravel Switch, KY: The Control Systems Group, 1989, 300 pages. A collection of previously published papers.

William T. Powers, LIVING CONTROL SYSTEMS II: Selected Papers. Gravel Switch, KY: The Control Systems Group, 1992, ??? pages. A collection of previously unpublished papers.

Marken, Richard S., editor. PURPOSEFUL BEHAVIOR: The Control Theory approach. Thousand Oaks, CA: Sage Publications: American Behavioral Scientist, special issue. Vol. 34, Number 1. September/October 1990. 11 articles, 16 contributors, 121 pages. A very readable introduction to a science of purpose and supportive research. Highly recommended.

Runkel, Philip J., CASTING NETS AND TESTING SPECIMENS. New York: Praeger, 1990, 186 pages. Contrasting the proper and improper uses of statistics with modeling for understanding and prediction of people as well as processes. Highly relevant to TQM efforts!

Hershberger, Wayne, editor, VOLITIONAL ACTION, CONATION AND CONTROL. Advances in Psychology 62. NY: North-Holland, 1989. 25 chapters, 33 contributors, 572 pages.

Ford, Edward E., FREEDOM FROM STRESS. Scottsdale AZ: Brandt Publishing, 1989, 184 pages. A highly readable introduction to PCT and a personal problem solving guide. The most accessible text available. Written as a comprehensive counseling story anyone can relate to.

Gibbons, Hugh, THE DEATH OF JEFFREY STAPLETON; Exploring the Way Lawyers Think. Concord NH: Franklin Pierce Law Center, 1990, 197 pages. Textbook for law students which spells out how lawyers think by explaining and using a PCT framework.

McClelland, Kent, PERCEPTUAL CONTROL AND SOCIOLOGICAL THEORY. 1991. This unpublished paper suggests that individual control as a phenomenon is central to understanding sociology. <st>

McPhail, Clark, THE MYTH OF THE MADDING CROWD. Hawthorne, NY: Aldine de Gruyter, 1990, 265 pages. Explains group behavior as a function of purposeful individuals.

Petrie, Hugh G., DILEMMA OF ENQUIRY AND LEARNING. Univ. of Chicago press, 1981. Discusses learning with explicit recognition of PCT insight.

Richardson, George P., FEEDBACK THOUGHT IN SOCIAL SCIENCE AND SYSTEMS THEORY. Univ. of Pennsylvania Press, 1991, 374 pages. A review of systems thinking in history, cybernetics, servo mechanisms and social sciences. Provides a perspective placing PCT in context in relation to other paradigms of human behavior.

--Dag Forssell with Gary Cziko 920611

Date: Fri Oct 02, 1992 12:49 pm PST
Subject: MIT talk

[From: Bruce Nevin (Fri 921002 13:14:03)]

I'm back from my month's field trip among the Indians, jet-lagged, and with much to catch up on. The following of possible interest was forwarded to me from the MIT AI list.

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

WHY DO WE SEE THREE-DIMENSIONAL OBJECTS?

Thomas Marill

October 1, 1992
4pm
8th floor playroom

Abstract

When we look at certain line-drawings, we see three-dimensional objects. The question is why; why not just see two-dimensional images?

In this talk we explore the theory that we see objects rather than images because the objects are, in a certain mathematical sense, less complex than the images.

However, there are infinitely many objects that project to any drawing. As a result, a second question arises: Given that we are going to see an object when we look at a drawing, which one will it be?

The theory under discussion holds that the object selected by the vision system will be the least complex of the available alternatives. Experimental data supporting the theory will be reported.

This work is based on the pioneering ideas of Solomonoff and Kolmogorov, and on the more recent ``minimum description length'' concept of Rissanen.

****Revolving Seminar****Revolving Seminar****Revolving Seminar****

The schedule for the rest of the semester:
Oct. 8: Bart Selman
Oct. 14: David Liddle
Oct. 22: open
Oct. 29: Eric Sven Ristad
Nov. 5, 12, 19: open
Dec. 3: Brian Subirana
Dec. 10: Andrew W. Moore

If you are interested in giving a talk send email to mdlm@ai.mit.edu. The Revolving Seminar has a small budget for reimbursing the travel expenses of senior researchers. If you are interested in a particular speaker, please let us know. We are particularly interested in inviting people who espouse views that are not widely represented within the lab.

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

Date: Fri Oct 02, 1992 3:08 pm PST
Subject: Re: Progress, Diagrams

Dag,
Actually, my interchanges with Bill were public, so if anyone else is keeping archives, they will be there.

The hypercard stacks are too big for Bill Silvert to keep. CSG-L generates about 1 Meg per month. I don't know how to make them generally available. They could certainly be uploaded to any ftp site that would give me permission to do so.

Martin

Date: Fri Oct 02, 1992 3:41 pm PST
Subject: Testing

[From Rick Marken (921002.1600)]

The following is a test of the emergency PCT understanding system. This is only a test:

"PCT makes all current textbooks on psychological methods obsolete."

True or false.

Have a nice weekend.

Rick

Date: Fri Oct 02, 1992 3:44 pm PST
Subject: Re: Hypnosis

[Martin Taylor 921002 19:00]   (Clark McPhail 921001)

> Then I discovered with work of Theodore X. Barber, a
>psychologist who is now retired but who examined the phenomenon of hypnosis
>across a forty year span of time. Simply stated, Barber rejects the
>"trance state" theory of hypnosis and has generated a considerable body of
>empirical evidence supporting his critique. Further, Barber has advanced
>an alternative interpretation which turns the commonsense view of hypnosis
>on its head. Instead of the subject-as-passive receptacle of the
>hypnotist's suggestions, Barber construes subjects as exercising variable
>degrees of imagination; that is, they are variably capable of imagining
>the outcome the hypnotist "suggests" and then carrying out the actions
>require to fulfill what they have imagined.

I have no PCT approach to hypnosis, but I will support the claim that the skill to be hypnotized is in the subject, rather than the skill to hypnotize being in the hypnotist. Here's a reference: Hypnosis for the Seriously Curious, K. S. Bowers, Monterey: Brooks/Cole, 1976.

Around the time this book was published, I had a summer student whose graduate work was with Bowers on hypnosis, and I helped them with the analysis of some data. The task, as I remember it (somewhat vaguely over this remove in time), was for the subject to shadow vocal material presented through an earphone to one ear, and to push a button when some target pattern appeared in vocal material presented to the other ear. Of course there were control conditions when only one task was to be done, and all of it was done with highly skilled hypnotic subjects under hypnosis or not. Without going into the forgotten details of the analysis, the result was that without hypnosis the tasks interfered with each other, but under hypnosis they did not. It was as if the hypnosis could modularize the whole perceptual feedback loop so that they did not use any common paths (in PCT terms, as if there were, for that purpose, two separable hierarchies that in a normal state were in some conflict).

This experiment convinced me that hypnosis is a real phenomenon. The lack of conflict does tie in with the notion of "variable degrees of imagination," since conflicts can be omitted in imagination even though they would arise in "real life." In the experiment, there was no intrinsic reason for conflict between the two loops, except possibly that both ears could hear the subject's own voice. But in the non-hypnotic situation, there was conflict, perhaps because in the usual real world one does not get indpendent information being presented to the two ears.

Speculation, guesswork, non-science. But fun.

Martin

Date: Fri Oct 02, 1992 4:54 pm PST
From: William T. Powers
          MBX: POWERS_W%FLC@vaxf.colorado.edu
TO: * Dag Forssell / MCI ID: 474-2580
Subject: figures

Dag, I found two copies of the figures for Arm Version 1. They will go into the mail for Saturday pickup.

Bill P.

Date: Fri Oct 02, 1992 7:05 pm PST
Subject: Re: Misc subjects

[Martin Taylor 921002 22:30]   (Bill Powers 921002.0600)

> But the idea of "competing behaviors" inhibiting each
>other doesn't seem very plausible to me. How does one behaving
>subsystem know which other behaviors to inhibit and which to leave
>alone? I don't think you could make a runnable model out of this, but
>you're welcome to prove me wrong. Your idea isn't even self-

>consistent, because you follow the above by saying
>
>>It is the world that stops them all happening at once. If they
>>could, they would.
>
>But that means that they are NOT mutually inhibiting each other inside
>the system. They're all in a state of perpetual conflict. This seems
>like a lousy design for a control-system hierarchy where you would
>like the loop gains all to be quite high.

Quite right. I never intended to imply that the competing behaviours had any connection internal to the hierarchy other than that they all had references to which a higher-level ECS contributed.

They certainly should not inhibit one another internal to the hierarchy, unless we are getting into the realm of pre-planning, which was the kind of choice behaviour I was trying to show was not always required (if ever).

I see you still don't buy the notion that not all controllable percepts are at any one moment controlled, and that most ECS gains either are very low or the ECS is in some way disconnected from the physical world. It still seems to me inevitable that there exists a very small actively controlling subset of the ECSs at any one level at any moment, the rest not, at that moment being in active control of anything except their imaginations. The idea of "giving up" is inherent in the normal situation.

>>I guess in the background there is another point I meant to discuss
>>at some time--"giving up." When there is persistent error, one
>>possible and often used response of an ECS is simply to reduce its
>>gain to zero, to give up on a hopeless situation.
>
>I have already brought up this idea on the net, but not by supposing
>that every ECS can perceive and judge "situations" as being
>"hopeless." My solution was a comparator that has a curve that I drew
>like this:

```
>
>     ^                  *      |
>     |                *    *   |
>Error|signal      *         *  |
>            *                * |
>*                   -         *|            +
>-------------------------------*-----------------------------------
>         <-- actual error-->   |*                          *
>                               | *                *
>                               |  *          *
>                               |   *     *
>                               |      *
>
```

>In the central region, feedback is negative. When a disturbance gets
>so large that the peak of the curve is reached (in either direction),
>the error signal begins to fall for further increases in disturbance.
>This results in a drop of output, still more error, still less error
>signal, and so on to the "giving up" regions at the ends.

This is a different situation, in itself quite legitimate. In the situation I am considering, the error need not be desperate, but it is consistently in the same direction. If the gain depends on the integrated error, I guess your curve would deal with the same situation, but the way I see it, the evolved solution should be expected to contain a switch. It would, as you say, be a lousy way to build a control hierarchy if every unachievable reference led to a continuous "kicking against the pricks." And in your system, it looks as if what would happen in the car-bicycle-walk scenario is that when the destination was near (and the higher reference-percept difference small) the gains of the unselected transport-medium-percept controllers would increase, giving rise to renewed conflict.

> It seems simpler to me to invoke a higher-level system
>that selects one of several means of locomotion (the kind of
>explanation you used for explaining why the wastebasket wasn't moved
>on the third day). This has the advantage of not pitting the outputs

>of high-gain (competent) control systems against each other, which
>always leads to a lessening of the capacity to control if not its
>complete destruction. I don't think you're taking full advantage of
>the concept of hierarchical control, at least not consistently.

I thought I was pointing out another instance of the power of hierarchic control. I am perfectly willing to believe that I'm not taking full advantage of its power, since I keep finding out new aspects of that power, and I doubt that discovery process has come to an end. But putting the onus of selection onto a combination of the relative insistences of the "competing" ECSs and the impedances of the relevant (mutually inhibitory) CEVs seems to me to be consistent with a lot of feelings one has when actually making that decision, besides not requiring any special mechanism of choice in the higher-level ECS. If the bike is handy and the distance not too far, one might use it, whereas if the bike were hidden behind piles of junk that had to be moved, one might not. On the other side of the person-world interface, if there were a reference level for keeping fit and using the bike would reduce that error as well as the error of not being at the destination place, then one might take the bike even if it were hidden by a pile of junk.

Perhaps it is worth noting that the concept of "impedance" applies not only to the CEVs of the world, but also to the reference-percept relation of a single ECS as seen from a higher level ECS that contributes to its reference signal and receives its perceptual signal. That impedance will be very low if the lower ECS has a high gain and the CEV it is controlling has a low impedance. And it will be near zero if the lower ECS is imagining. Conflict at the lower level may well lead to the higher ECS seeing the impedance of the lower as being rather high. And for good measure, one should mention the possibility of negative impedance (which implies positive feedback, and trouble).

The world may well be much more closely coupled than the perceptual control hierarchy is. I think that may be what people mean by "leaning on the world."

Martin

Date: Sat Oct 03, 1992 3:59 am PST
Subject: Sundry subjects

From Greg Williams (921003)

Gary Cziko: I received log9209c yesterday (from "CSGnet God," no less!) BUT NOTHING ELSE.

>Dag Forssell (921001-1)

>In the meantime,
>can you make any suggestions on the arm demo and crowd using my format?

Just that to the extent you can somehow have the audience INTERACT with them, rather than simply observe them, I think you will see more "aha" actions.

>Clark McPhail (921001)

>My
>impression for years was that hypnosis works because the subject focuses
>upon and tells him/herself to do exclusively what he/she is instructed to
>do by the hypnotist. Then I discovered with work of Theodore X. Barber, a
>psychologist who is now retired but who examined the phenomenon of hypnosis
>across a forty year span of time. Simply stated, Barber rejects the
>"trance state" theory of hypnosis and has generated a considerable body of
>empirical evidence supporting his critique. Further, Barber has advanced
>an alternative interpretation which turns the commonsense view of hypnosis
>on its head. Instead of the subject-as-passive receptacle of the
>hypnotist's suggestions, Barber construes subjects as exercising variable

>degrees of imagination; that is, they are variably capable of imagining
>the outcome the hypnotist "suggests" and then carrying out the actions
>require to fulfill what they have imagined.

A reference I recently found which (citing Barber extensively, but many other investigators also) gets into the "active subject" business -- which makes just as good PCT-sense in hypnosis as in any other kind of situation where one person is controlling for seeing certain kinds of actions by another person is Graham F. Wagstaff, HYPNOSIS, COMPLIANCE AND BELIEF, St. Martin's Press, New York, 1981. Its chapters include "Hypnosis Past and Present," "Sham Behaviour and Compliance," "Compliance and Hypnosis," "How Do I Know I'm 'Hypnotised'?" "Some Hypnotic 'Feats'," "Further Characteristics of Hypnotic Performance," "Differences in Hypnotic Suggestibility," "Hypnosis and Pain," "Hypnotherapy," and "The Nature of Hypnosis."

I was once (maybe) hypnotized by shrink/neurophysiologist Jerry Lettvin (at MIT; best known for the "What the Fly's Eye Tells the Frog's Brain" paper, which he co-wrote with Warren McCulloch and Humberto Maturana, among others). Whatever was going on, to me it seemed most akin to deliberate (on my part) play-acting. I was trying to play the role which Jerry, as "director" wanted me to play, because, as I recall, I wanted to humor Jerry and thought it harmless. I never felt as if I weren't "in control."

>Gary Cziko 921002.0200 GMT]

>Relevant to the recent discussion of purposeful influence and
>reorganization is an excerpt (sans emphases and footnotes) from a chapter
>of my book in preparation.
>...
>I currently feel that it is indeed possible for people (e.g., teachers) to
>influence long-lasting changes in others' (e.g., students') control
>systems.

Ditto. For me, the question isn't whether it is possible at all, but to what extent, under which conditions.

>We can also easily imagine that the learner in this example would be very
>highly motivated since failure to learn could result in death. From a
>perceptual-control-theory perspective, motivation simply refers to error
>(that is, a difference between a perception and the reference level for
>that perception) which results in action to eliminate the error (see Figure
>7.1). From this perspective, motivation is considered to be internal to
>the student since the reference level of the controlled variable is
>determined by the student, not by the environment.

If "motivation" = difference between perception and reference level, then motivation is NOT SOLELY determined by the student, but jointly by the student and her environment (since perception is affected by environmental disturbances independent of the student). Nevertheless, since the comparator is internal, the error SIGNAL is "internal to the student."

>Models and
>instruction can provide useful information in the form of constraints of
>what not to try, but they cannot provide explicit instructions concerning
>exactly what to do.

Well said. Models and instruction can only influence (or "guide"), not determine, what the student does.

>In addition to allowing more time for the learning to take place and
>providing constraints in the form of models and verbal instruction, the
>teacher can also provide easier access to the knowledge or skill by
>providing a series of less demanding intermediate goals. One way is to
>break down the skill into a number of subskills and provide opportunities
>for the subskills be acquired.

Yes, and this could make the "guiding" quite efficient, I claim (contra Bill).

>By breaking down a complex problem into easier subproblems learning is
>facilitated since the probability of finding a solution to each subproblem
>is higher than that of finding a solution to the more complex
>problem--success in learning to make effective arm alone in swimming is
>more likely than success in learning to make both arm and leg movements
>together.

Here I think you need to talk about why the probability of finding a solution to a subproblem is (usually) easier than finding a solution to the whole problem in one swell foop. My hypothesis is that the trajectory from control system state A (initially, before learning) to state B (after subproblem is solved) is "shorter" in some abstract space sense than the trajectory from state A to state Z (after the whole problem is solved). This could stand considerable fleshing-out.

>In other words, the teacher
>arranges the environment so that the student is continually encountering
>error, but error that is not too large so that the student's reorganizing
>efforts are likely to be successful and set the stage for the next
>introduction of error.

Sounds reasonable to me (for whatever that's worth!).

>Bill Powers (921002.0600)

>Greg Williams (920928) --

>I agree with

>>1. A disturbs particular perceptions being controlled by B so that
>>B compensates for the disturbances with actions which A wants to
>>perceive.

>>2. A arranges B's environment so that when B controls for
>>particular perceptions, A perceives what he/she wants to perceive.

Well, that says a LOT. Should keep some sociologists busy for a while!

>but I have a problem with

>>3. A arranges B's environment so as to trigger learning
>>/reorganization in B's control system resulting in actions which A
>>wants to perceive.

>>4. A applies physical constraints or threatens to apply physical
>>constraints to B so that B's actions are as A wants to perceive.

I would have predicted your having a problem with 3, on which we have previously agreed to disagree, but not with 4, which (perhaps in confusing form which threw you off?) is just "controlling another by use of overwhelming physical force or threat thereof."

>According to my model of what triggers reorganization, these would
>both mean arranging the environment so that B suffers critical error
>(you notice my return to Ashby's term) such as hunger, thirst, pain,
>illness, suffocation, "stimulus deprivation," or whatever you want to
>put on the list.

Note that I have included BOTH "learning" AND "reorganization" in 3. Do you think we DON'T need to postulate something "less" than reorganization in cases which Gary has raised, such as learning how to multiply (where it appears that (1) the learner's control system is altered and (2) there is no critical error triggering the learning)?

>By definition, reorganization is unsystematic.

Its TRAJECTORY is unpredictable, but its successful endpoint isn't necessarily unpredictable -- though it might be unpredictable sometimes in practice, depending on the particulars of the situation. Of course, it might require infinite time to get to the predicted endpoint.

>This means that you can't
>predict what behavior will be used to correct the error unless you have
>removed all means of correcting it but one, which is within B's capacity to
>learn.

That's basically what I (and Gary, I think) have been contending. To predict the endpoint requires providing certain constraints so that the problem can only be solved in a finite number of ways.

>That's easy to do with a lower animal or a child, but hard to do with
>an adult human being.

Not necessarily in any case. Verbal interactions with typical adults could make "guiding" their learning much easier than guiding the learning of a non- human organism or a pre-verbal child. But I am sympathetic to the implication that adults typically have more POSSIBLE trajectories for a learning process.

>You do note that these methods involve conflict, but you
>don't mention that the outcome is largely unpredictable because
>reorganization is involved.

"Largely" is, of course, in the eye of the beholder. I look around and see people becoming capable of solving many kinds of problems set for them by others -- with (sets of) actions "largely" predictable by the problem-setters. (I say "sets of" because, often, the problem-setter doesn't care which EXACT actions are used to solve a problem; as long as the PARTICULAR actions are part of the set of all-actions-which-solve-the-problem, that's fine with the problem-setter -- that's what the problem-setter is controlling for seeing. In fact, the problem-setter might not know ANY solution to the problem in advance, or how to solve it himself/herself.)

>In the second section, I don't understand...

I'll address these questions in a later post, when I have more time.

>Tom Bourbon [921002 -- 10:50 CDT]

>It is not that the model "does good" and fits only
>if people "do good." Rather, the model does as well or poorly ad
>the person.

A perfect tracking model would exactly replicate what the person does, in terms of cursor movement. I just meant to say that I believe the high (not exact) correlation between simple PCT tracking models and what people do is attributable largely to the fact that, for much of the time in a given run, control is "good." If a run were made which consisted mainly of sudden "jerky" disturbances, then, for a high-correlation model would (I hypothesize) require complications similar to those in the VERY complex tracking models developed by human factors engineers.

>Now we see some model-person correlations below .99, but they still
>exceed by orders of magnitude the modal and median correlations in
>the behavioral literature.

I encourge you not to give up the PCT goal of 0.99+ correlations. Otherwise, I believe, you will not be able to decide WHICH PCT model is the correct one. Fairly low (but higher than typically found in the behavioral sciences) correlations are convincing (at least to me) that PCT models are correct, generically, for tracking. But I next want to know WHICH PCT models are better than others, and this requires comparing various PCT models' correlations with what people do UNDER CONDITIONS AFFORDING ENOUGH SENSITIVITY TO SHOW SIGNIFICANT DIFFERENCES IN CORRELATIONS. Such conditions, I claim, are "poor" control conditions.

>I like your emphases that indicate the relativity of "good and bad"
>control and of "high and low" performance. Your reasoned treatment

>is better informed than that of people who merely assert their beliefs
>and prejudices on these matters. PCT models control. When control
>is present, PCT predicts with precision -- even if the control is
>mediocre. It predicts with great precision if the control is precise.

Yes. Now it's time to begin choosing AMONG models which are based on PCT.

>Bruce Nevin (Fri 921002 13:14:03)

>If you are interested in giving a talk send email to mdlm@ai.mit.edu.
>The Revolving Seminar has a small budget for reimbursing the travel
>expenses of senior researchers. If you are interested in a particular
>speaker, please let us know. We are particularly interested in
>inviting people who espouse views that are not widely represented
>within the lab.

I move and second that we propose Bill Powers as a speaker for MIT's Revolving Seminar. Maybe we could send a video camera along -- I'd love to see Bill and R. Brooks interacting (or NOT interacting, as the case might be) face-to- face!!!

Off to Cam-nirvana (Cam's my 13-year-old son): the annual Lexington stamp show. Anybody out there got a Mint Never Hinged Sweden #1710 (spider)?

Best wishes,   Greg

Date: Sat Oct 03, 1992 5:43 am PST
Subject: the perceptual basis of HPCT

[From Bill Powers (921003.0600)]

Greg Williams (921002) (re-posted by Dag Forssell, thanks) --

>Quite frankly -- I've told Bill, of course -- I have never bought
>into Bill's hierarchy AT ALL except as a (not very well fleshed-
>out) existence proof for AN instantiation of PCT details. At this
>stage, I think PCT has a lot to learn from living organisms before
>pushing HPCT as THE instantiation.

It's interesting how easy it is to overlook the one living organism from which we can learn the most, and the most directly: ourselves. We experience directly the perceptual activities of our own brains. Our brains deliver to us a world whose organization reflects the types of perceptions the brain is equipped to derive from its raw inputs. By examining that world as evidence about the organization of perception, instead of as evidence about a physical outside universe, we can find regularities and dependencies in it that are universal, independent of any particular scene or experience yet common to all. This is the basic principle behind HPCT.

Unfortunately we become accustomed from earliest childhood to accepting this world as an external given, as something that exists independently of us while we peer at it through two windows looking out of the dark cave inside our heads. All the action is outside the cave; that's where everything is and where everything happens. We stand permanently in the mouth of the cave, from which we can see approximately a hemispherical view of that world; we can reach out and touch it, we can smell it and taste it and exert efforts on it and bring pieces of it inside us (where they disappear).

The result is that when we talk about studying perception, we look in the wrong place. We try to turn around and look into the depths of the cave, turning our backs on that bright and solid world outside it. But of course there are no perceptions inside the cave. We can't even turn around to look in there. The best we can do is to shut our eyes or wait for dark and imagine another world. So the first thing we do in trying to understand perception is to shut it off and ignore it, and create an imagined world -- still outside us, the observers, of course -- with which we try to explain perception.

We try to imagine what is going on inside someone else's brain. We open up other brains and put electrodes in them, and look at the displays on oscillographs where trains of blips go marching by (all the time perceiving these things). We believe that we have learned something valuable about perception, because now we can look at the blips, and look at the world impinging on the other person's senses, and see a correlation between the blips and the world. We tell ourselves that behind our backs, as we stand at the mouths of our personal caves observing and doing these experiments, there is another gray mass of filaments like the one before us, and that inside our caves, behind us where we can't see, there is a universe of blips like those on the chart or display screen outside us where we can see it. In this way we get the feeling of knowing a little about our own perceptions.

All the while, of course, perception is displayed before us with a resolution that reveals the tiniest details of the smallest objects, completely filling the field of vision and taste and smell and feel and sound. And we ask ourselves, "How can I find a way to study the brain's perceptions of this world?"

When the physical sciences began to develop, we discovered a strange thing about this world outside us: it is full of things we can't see, taste, smell, feel, or hear. It contains mass, heat capacity, charge, density, atoms, compounds, energy fluxes some of which pass invisibly right through us, attractive and repulsive forces, and vacuums. This discovery led to great confusion: how can it be that we experience one world directly, and a completely different one through our scientific instruments? What is this world that we see, hear, and so forth, if it is not the physical world itself? Most scientists solved this problem by ignoring it; they accepted the instrument readings as the true picture of the world, the physical-chemical world, and dismissed the world of common experience as an incompetent misrepresentation related only loosely and ambiguously to the REAL world. Just why this incomplete and unreliable world should exist outside us, and why our senses should respond to it rather than to the real world, nobody could explain. The world to which our senses respond (and that is how they put it) was, at best, an illusion. Studying it had no possible scientific value.

So the only valid way of studying perception became that of comparing the readings of scientific instruments against the readings of other scientific instruments attached to the brain, or of studying scientifically measurable responses of the muscles when physical and chemical events occurred. The world of direct experience, having no scientific value, was left to laymen as an illusion to keep them busy.

There is, of course, another interpretation of all this, the one underlying HPCT. It is that the world we experience directly has ALREADY passed through the sensory inputs by the time we experience it. It is the direct manifestation of the brain's way of reading the same world that scientific instruments read. It is the brain's way of realizing, of making apparent, the ordering of the universe. Scientific instruments create a different realization of, we presume, the same underlying order.

If the direct and the scientific views of the world give different information, that is because the brain's realization is organized differently from the devices that produce scientific meter readings. To understand how the brain's perceptual system is organized, therefore, we should look for aspects of the experienced world that are not represented in the world of scientific meter readings. We should look at experience as consisting of a different sort of meter readings (in which, of course, scientific meter readings themselves are embedded). We should look for properties of the directly experienced world that are not among the properties we derive from observation through scientific instruments.

These properties are not hard to find, once you realize that they are properties of perception. There are, for example, relationships. One object can be larger or rounder than another, or above or beside another. Between the objects there seems to exist a connection, a comparison, a dependency in which each of the objects adds meaning to the other, or to both. Having seen a relationship like the distance between two objects, we can then derive a scientific meter reading from it, by placing a meter stick to span the objects and reading off numbers from it in which to quantify distance. But in all of physics, there is no THING called the distance between the objects. Between the objects there is nothing physical that belongs to the objects in the same way that their separation in space belongs to them in direct experience. There is no physical quantity called larger-than, or rounder-than, or above, or beside.

Once our attention is called to this "added value" in direct perception, we notice other kinds. We notice, for example, the objects that are related. Each object is a recognizeable thing, a shape in two or three dimensions. There is no physical quantity called "shape," although perceived shapes can be approximated by physical entities made of spheres, cones, rhomboids, and the like, and can be measured through recording large numbers of meter readings. Shape is one of the brain's ways of realizing the order in the universe.

As we get used to looking for non-physical properties of the experienced world, more and more classes of properties appear. We see that motion is a non-physical property; that the elements of vision such as color and shading and so on are non-physical properties. Each such property can be associated with meter-readings, given suitable instruments, but the meter readings deliver up only numbers. There is no meter that can say THIS is a shape, but THAT is a rotation. Such realizations are the exclusive purview of the brain.

Continuing this exploration, we eventually pass out of the regions of experience where scientific instruments can provide any readings. We find, for example, that the world is divided into classes of things. There is no meter that can indicate the class to which some arbitrary item belongs -- indeed, one item might belong to many classes. We find that sequence or ordering exists; that ABC is not the same as CAB. While our brains can perceive the ordering in meter readings, no one meter reading indicates an ordering in the physical world. Ordering is not a property of the physical world.

We see, too, that the world is logical; that A can cause B, in the sense that it is NOT true that A occurs and B does NOT occur, although B may occur independently of A. We discover that at least one kind of causation in the world is identical in form to logical implication. This entity, causation or implication, has not even the remotest existence in the world represented in scientific meter readings. It is the brain's realization of another kind of ordering in the universe: a perception of a higher level than the perception of A and B.

The longer we continue this kind of exploration of the world of direct experience, the more personal become the discoveries; we place the very essence of rationality into the world of perception created by the brain. And that is the point of HPCT.

The point of HPCT is to notice everything that exists in the world of direct experience, whether we class it as "inside" or "outside," "concrete" or "abstract," and place it into a model of the organization of perception, which is, we say, a model of one aspect of the brain's functions. This is not a "scientific" endeavor; it can't be, because the meter readings which are the basis of physical realizations of the world are themselves simply elements of the world of direct experience. This is a process of noticing kinds of perceptions which are not part of the physical realization of the world, which "add value" to physical meter readings, and of putting the components of experience thus recognized into an orderly system of understanding, which we call a model of the brain. And in doing so, we notice, of course, that we are doing this, and we add to the model a kind of perception we can call system concepts.

With all these aspects of experience thus separated into discernible kinds, we can ask how they are related. Here and there we can see certain dependencies. For example, there is never an object but that it is composed of sensations. We think of the sensations as attributes of an object, but it is also true that when all attributes of an object are removed, there is no object left, save in imagination. So we see that objectness somehow depends on the attributes of objectness, which we label sensations. But we can imagine an object even in the absence of color, edges, and so on, so we suspect that the perception of objectness is somehow derived from and exists separately from the individual perceptions of sensations. And in general, we can see that the attributes of any perception are themselves perceptions, with the attributes acting like the arguments of a function, and the perception "having" those attributes acting like the value of the function. From there it is only a short step to a neural model in which all perceptions are neural signals, related by computing functions that make one signal dependent on the value of signals at a lower level.

The current form of the HPCT model describes what are purported to be eleven kinds of perceptual realizations that we see in the world as we stand in the mouths of our caves. We see all these aspects of perception not behind us, in the cave, but in front of us, out there in that unscientific world of direct experience. Even such nonmaterial things as relationships and sequences and logical propositions and principles and system concepts we see out there, as aspects of that world. Only when we imagine does that world appear to move inside -- but then we are attending neither outside nor inside, but in a different way.

Each person must discover the components of perception and their relationships to one another, by careful examination of direct experience. Only when each person does this, independently, is it possible to reach a consensus on whatever aspects of direct experience are in fact common to all human beings. This is as public as any observation can be, even though the basic idea says that all observation is private. We reach consensus through communication, through asking if you're experiencing what I'm experiencing and if you call it what I call it, and if it behaves for you as it behaves for me in circumstances we classify in the same terms. And of course, under HPCT, we control these perceptions, and ask others to control them also, and look for conflicts that arise from differences in our perceptual worlds. In this way our worlds become mutually consistent, although we can never know whether they are identical.

If fewer people took the world of direct experience for granted, and if more people began to see it as evidence of how the brain perceives (and controls), the organization of HPCT might seem less arbitrary, less like a pronouncement handed down from on high and more like a description of one's own world.

Best,   Bill P.

Date: Sat Oct 03, 1992 6:34 pm PST
Subject: Re: Sundry subjects

[Martin Taylor 921003 12:15]   (Greg Williams 921003)

>>=Tom Bourbon
>>Now we see some model-person correlations below .99, but they still
>>exceed by orders of magnitude the modal and median correlations in
>>the behavioral literature.
>
>I encourge you not to give up the PCT goal of 0.99+ correlations. Otherwise, I
>believe, you will not be able to decide WHICH PCT model is the correct one.
>Fairly low (but higher than typically found in the behavioral sciences)
>correlations are convincing (at least to me) that PCT models are correct,
>generically, for tracking. But I next want to know WHICH PCT models are better
>than others, and this requires comparing various PCT models' correlations with
>what people do UNDER CONDITIONS AFFORDING ENOUGH SENSITIVITY TO SHOW
>SIGNIFICANT DIFFERENCES IN CORRELATIONS. Such conditions, I claim, are "poor"
>control conditions.

Please don't fall into the trap of thinking that "significant differences" have any meaning other than that an experiment is sensitive enough to show them. They have no relevance to the real world, ever.

Good control, almost by definition, obscures the structure of the controller. Finding high correlations means two things: (1) Once again, here is a situation in which PCT works [would that be news, except to someone with a belief that it wouldn't?], and (2) the experimenter intuited a CEV pretty close to the one being controlled, and that control was being performed with high gain.

If one is interested in finding out HOW the control is being done, that information is mostly to be found in the 1% error that remains. It is easier, but less interesting, to find structure in the error when control is poor. It is less interesting because there are several possible reasons why the control may seem to be poor, one of them being that the experimenter has not intuited a CEV close to the one really being controlled.

I think it is the very high correlations achieved by PCT models in various situations that provide the opportunity really to tease out what is happening in the brain. S-R approaches do it the opposite way, by (partially) breaking the feedback loop. That prevents the subjects from behaving in a way that they can control, but provides a wide latitude for variations that indicate structures or functions internal to the mind/brain system. There are too many sources of variability to make it easy to find which components are important, and there are therefore many competing schools of thought on what is going on. Within PCT, there is a much better opportunity to concentrate on the effects that are really due to structural/functional factors.

Let's consider an example brought up by Rick Marken (920929.1000):

> For example, in my "area" vs
>"perimeter" control study, I did the test for the controlled variable to
>determine whether the subject is controlling x+y vs x*y (where x and y
>are height and width of a quadrilateral figure). Using x+y as the
>hypothesized controlled variable the error in predicting responses was,
>I think, about 2%. With x*y as the hypothesized CEV, the error was halved,
>to 1%. You could probably do slightly better with some other hypothesized
>CEV -- maybe sqrt(x*x+y*y) -- but clearly you are on the right track
>with x*y.

The interesting evidence is not that people control with high accuracy, but that the residual error is different if one presumes different perceptual functions (and therefore CEVs). If it turns out that x*y is the best that can be achieved, the residual error probably includes the failure of lower-level control systems to achieve their reference levels. But it also signals the _effective_ gain of the x*y control loop, and much more important it shows that something in the system is capable either of computing a product or of adding logarithms. Which is it? The statistics can give a clue. If the residual variance is proportional to x and to y, logarithms are the likely answer. If it is more or less independent of x+y, then mutiplication is probable. What then? If a logarithm can be developed in one part of the hierarchy, is it not likely that it can be done in another part? Then perhaps we should look for logarithmic relations elsewhere in the hierarchy. But if it looks more as if the answer is multiplication, then a host of different relations seem reasonable candidates as the CEVs in other situations.

It seems to me that the high correlations support the idea that "PCT triumphs again," whereas the small errors, properly analyzed statistically, show how the brain works.

Martin

Date: Sat Oct 03, 1992 10:55 pm PST
Subject: reality of rules paper

For anyone interested in such things, I have a paper draft on the psych. reality of (some) linguistic rules, sort of a limited brief for cognitivism. Definitely no PCT in, but the stuff I talk about will have to be faced by a PCT theory of language.

Avery.Andrews@anu.edu.au

Date: Sun Oct 04, 1992 12:55 pm PST
Subject: Modelling bad control;learning swimming; control via reorg

[From Bill Powers (921004.0800)] Tom Bourbon (921002.1050) --

Your treatise on good control vs good modeling makes a critically important point, which is that the control model doesn't fit behavior well just because it's a control-system model. It has adjustable parameters that can make it fit both excellent control behavior and very poor control behavior. With different values of the parameters, it would fit even expert control behavior very poorly.

The full significance of the numbers you posted isn't self-evident in the correlations:

A	
c-t	mod-pers
.987	.984
.987	.984
.463	.880
.404	.954
.430	.930
.986	.989

The "c-t" column indicates, via correlations, control error by the real person. In three of the cases, the person's actual "cursor position" (value of the variable A) departed by a very large amount from the pattern it should have followed for perfect control ("target"). With a correlation of only 0.404, explaining the observed behavior simply by saying "it's a [perfect] control system" would have accounted for only 16% of the variance of variable A even using a measured reference level. But for that particular entry, explaining the behavior by saying "It's a control system with a [particular measured] value of integration factor and a [particular measured] reference level" accounts for 91% of the variance of variable A. Even for the entry in which the model did the most poorly and the subject did the best (among those worst three), the percent variance accounted for rises from 21.4% to 77.4% when the parameters of the model are adjusted to fit the behavior of the person, rather than for perfect control.

One lesson to be learned from this is that it means little to explain behavior by saying "it's a control system." That explanation could account for 90 percent of the variance of observed behavior (or better), or for 16 percent of the variance (or worse). It's not just the architecture of a control system that enables it to predict behavior; it's the quantitative parameters. When we say that a control-system model explains behavior, we mean a control-system model that has the right parameters, not just the bare boxes with arrows between them. The qualitative model applies only when control is expert, so there is little difference between the actual control processes and perfect ones. And when control is expert, you hardly need the quantitative model to prove that control is going on. -- Greg Williams (920928) --

>An aside: What's going on in HYPNOSIS in PCT terms???

Another possibility is "suggested" by the fact that hypnosis almost always involves verbal communication. You're getting sleepy, your arm is getting heavy, and so on. Induction processes seem aimed at drawing attention to lower-level perceptions while a voice drones on and on. If you can get the locus of awareness moved into the lower control systems, this leaves the higher ones running on automatic. Then words (and meanings) circulating through the logical and higher levels would be indistinguishable from self-generated logical processes or programs, and might begin to be used as if they were the person's own intentions at the higher levels, particularly the category through program levels. Disturbing principles or system concepts would lead to resistance, just as they say. And of course awareness is still present -- you feel as "in charge" as ever, except that what you happen to feel like doing is partly under external direction. Oh, blah, blah, blah.

Gary Cziko (921002.0200!) --

I like that excerpt from your book very much. It suggests that there is some conscious participation in and direction of reorganization -- at least direction of WHERE it is applied, if not direction of the detailed process itself.

As I was reading your examples, I was recalling my own first lessons in swimming. I really wanted to know how, but the main memory that persists is the terror involved in making myself go into the deeper water in order to be taught. I did not, after all, know how to swim at that point. The biggest problem was in keeping the fear (and excitement, pretty much the same thing) within bounds so I didn't just head for shallow water and give up, or panic and start thrashing around.

I think, as a theorist, that it was the fear/excitement that drove the reorganization, but clearly what I learned was swimming, not high- jumping. My existing goal structure, which included wanting to learn to swim like other people were doing, had to be involved, and the fear certainly resulted from my putting myself into a situation I felt as dangerous and wanted to leave (an action that could only have resulted from operation of my existing control systems and a more or less deliberate induction of conflict). Also, my mother and father were giving me verbal instructions and demonstrations, so clearly what I was learning was something presented to my senses in the present-time world (as opposed to a genetically-driven process). They not only taught, but reassured me of my safety because they were there to get me out of trouble, so this lessened the conflict.

One aspect of this learning process that's of interest is that it involved imitation of movements, using control systems already well- developed but not adjusted to the underwater environment. I was consciously trying to move my arms and legs in the demonstrated way, but DISCOVERING the sensory consequences of acting in just that way under water. I was taught pretty much as you outlined -- learning to tread water first, and stay afloat. But nobody could tell me how the resistance of the water was going to affect the way my limbs moved, or the way that the resulting forces kept me up. I could see how my mother moved her arms slowly back and forth, sideways, but only when I tried it did I realize how much force had to be used to create those slow movements. I did a lot of experimenting, discovering that moving a flat hand sideways would push me up if the hand was tilted correctly. I didn't consciously make any changes in my own organization; I just kept noticing effects and trying to repeat them (or avoid them, cough, choke). But clearly, the changes that were going on were all aimed at gaining control over the new perceptions that were being generated by these new ways of using familiar actions in a new environment -- repeating some perceptions and trying to make sure others didn't occur again.

I think we have to separate the process of reorganization itself, which is random but reasonably efficient, from the processes that create the need for reorganization and those that direct it to the appropriate systems. The hierarchy, as Greg Williams has been suggesting, does get into the act, in the sense that it can actually create a situation that induces critical error, and can thus (learn to) cause reorganization independently of any environmental pressures or accidents -- which, of course, continue to work through their own effects on critical variables. Mary, reading your chapter excerpt, pointed out the article in Science News about people who seek risk on purpose, like rock climbers. They put themselves in more and more risky situations, but by doing so acquire greater and greater skill. It must be confidence in their own capacities for reorganization that makes this attractive -- they know that they will quickly learn the required skills, so that the fear goes away. I suppose that they relabel their emotions, so that the sensations of bodily state that others might call fear are categorized in more positive ways, and not avoided.

This conscious creation of situations that turn reorganization on doesn't produce any control of reorganization itself; it merely creates errors that turn the process on, and confines it to the areas of the brain pertinent to learning a particular kind of thing. I don't think we're ever conscious of reorganization itself; only of a gradually increasing ability to make certain perceptions occur when we want them to, and of a heightened state that reflects the critical errors that drive reorganization. The human system, which begins with a very low degree of organization in comparison with most animals, can reorganize to the extent that it can even learn to create the conditions that produce reorganization. It still can't direct that process in detail, but it can assure that the process continues until the conscious goal is met.

Anyhow, your descriptions of learning to swim make it clearer than ever that even if my model of reorganization is basically OK, it is part of a larger system and doesn't tell the whole story.

Skinner, in his invention of "shaping," anticipated "scaffolding." He found that animals learn quickest when only small changes have to be made at any stage. This has important theoretical implications, I'm sure, although I couldn't say what they are right now. Somehow the space in which an organism gets from organization A to organization B must be continuous between A and B -- learning couldn't work in a universe where there was no coherent path from one mode of behavior to another, with intermediate steps also making some kind of sense.

Greg Williams (921003) --

I guess what I object to in your proposals that we can "guide" reorganization in other people is the assumption that the guider can arbitrarily set this as a goal and then just carry it out. The whole discussion of "influence," from your end, seems to be setting up the influencer as an independent agent, with the influencee simply reacting as the influencer wants. This puts all the volition and agency into the influencer and ignores that of the influencee. It really seems to me to be an attempt to get around the interactive nature of control and go back to the old picture of the organism's behavior as a passive consequence of external events.

When we influence the actions of others in systematic ways, we succeed (when we do) only BECAUSE they are independent controlling agents, not DESPITE that fact. You can't make a person reorganize without inducing critical error into that person, or creating a situation that does it for you. The immediate reaction of the other person will be to find SOME way to counteract this most invasive act, or to negate or escape from the situation that produces the error, by the most direct and immediate means possible including setting the situation back to the way it was. The other person isn't going to just wait around passively while you, to suit your own goals, start causing pain or suffering. You can set problems all you like, but there's no constraint on the other person to solve them according to the rules you had in mind, or to seek the particular solution you consider the "right" one, or even to be interested in the problem in the first place. When you use this method, you're in constant danger.

I objected that you can predict the outcome of reorganization only when you arrange the environment so that only one action or learned method of control will in fact correct the critical error. You took this to mean that there IS a way of guiding reorganization -- just do what I described. But I offered this as an objection because it is so unlikely that you could ever do this with a human adult, a peer, without some great source of power behind you. It implies that you know all possible ways of solving the problem to the other's satisfaction; it implies that you already have far more control of the local environment than the other person has, so you can freely rearrange the environment without any effective objection from the other person. It implies that you have sole control over what the other person needs.

I claim that in order to establish the conditions under which this kind of influence (and many other kinds) could be even apparently effective, you would have had to establish a background of power over the other person and over the environment, power denied to the other person. In many of your examples of means of influence, I believe you are assuming (perhaps without knowing it) surrounding conditions that amount to establishing a far greater degree of coercive control over the other person than the degree of influence indicated by the simple example. This is particularly true when you speak of influencing or guiding another person's reorganizations. But it is true even of such innocuous means of influence as rubber-banding. How are you going to keep the other from seeing what you're doing? Or, after seeing it, from objecting to it or deliberately changing the goal simply to frustrate your attempt at influence?

When we talk about influencing children in the classroom, we tend to overlook the fact that they are in that classroom, and are trying to do passably well in their courses, not out of their own choice but because they have been physically forced to attend school, stay in their seats, be quiet, listen, study, take tests, and keep trying until they pass the course. The children are not offered any choices at this level of description. It does not matter whether they want to be there, whether they want to learn the course material, whether they want to take the test. They MUST be there, and if they want to avoid pain and disgrace, they MUST pass the course. Throughout the years of mandatory education, the children are in prison. If they try to escape they will be seized, punished, and returned to class. The elementary and high-school system is fundamentally based on coercion and control by people who have the power to withhold rewards, restrict freedom, and punish. This philosophy naturally carries over into higher education. The situation is quite analogous to that of a rat in a Skinner box, where the experimenter can put the rat into the box regardless of where the rat wants to be, and give or withhold reward while restricting the means of getting it to performing just the action that the experimenter wants to see.

This is not "influence," but flat-out control. It is based on brute physical power capable of crushing any resistance to it.

The same sort of glossing-over of the background of brute power that we see in mandatory education also exists in the world of business. Even in enlightened business organizations where people are allowed to choose their own hours, their own tasks, their own ways of executing those tasks, and even some of the orders that management gives them, the underlying principle is that of control by overwhelming physical force. The very idea that people are "allowed" to do these things implies a surrounding wall of coercion. A person who does not want to perform any task at all will simply be left to starve. Any attempt to avoid starvation in a way other than achieving the company's objectives will be punished by removal of the paycheck; any attempt to get the paycheck anyway will be punished by the Law. No matter how varied and enriched the environment within the cage, it is still a cage in which the person is kept by brute force.

When, in the background, you have the power to control a person's physical well-being or even existence, the role of mild "influences" becomes an illusion. The company can offer "incentives" in the form of "bonuses" for good work, but the clear implication is that bad work will be punished by loss of the bonus; you can't have the bonus anyway just because you desperately need it, or because your salary has been cut to finance the incentive program. Mild, friendly, helpful influences succeed mostly because failure to be influenced in the direction obviously wanted will result in loss of the job. The objective of these influences, keeping the worker or student in conformity to the larger objectives of the owners of the business or the administrators of the school, is backed up by the credible threat of force -- credible because it is used whenever necessary, which is often.

You propose as one effective means of controlling another:

>3. A arranges B's environment so as to trigger learning
>/reorganization in B's control system resulting in actions which A
>wants to perceive.

The natural reaction of B is to prevent A from arranging B's environment in that way, perhaps by sinking a knife into A or by driving a car past A's house and letting loose with an Uzi. This method of control dooms itself to failure. That's why I don't treat it as a serious proposal: it can't work, except temporarily.

There is another reason I don't treat it seriously: that is because reorganization will NOT result in the actions which A wants to perceive, but in whatever actions B decides may work best for B. You assume that A can arrange the environment so that only ONE action can succeed. But you overlook the action of changing the environment back to its original state, you assume that it's POSSIBLE for A to make this arrangement without any interference from B, and you assume that A can anticipate all possible outcomes of reorganization (such as suicide). I think these assumptions are unwarranted, in general.

>I encourge you not to give up the PCT goal of 0.99+ correlations.
>Otherwise, I believe, you will not be able to decide WHICH PCT
>model is the correct one. Fairly low (but higher than typically
>found in the behavioral sciences) correlations are convincing (at
>least to me) that PCT models are correct, generically, for
>tracking. But I next want to know WHICH PCT models are better than
>others, and this requires comparing various PCT models'
>correlations with what people do UNDER CONDITIONS AFFORDING ENOUGH
>SENSITIVITY TO SHOW SIGNIFICANT DIFFERENCES IN CORRELATIONS. Such
>conditions, I claim, are "poor" control conditions.

I agree with you, mostly, but I think we should give up on correlations altogether. They're just not sensitive enough when they get up into the 90s. It's much simpler to use RMS error over peak-to- peak value of the predicted variable, which happens to be a standard measure of noise-to-signal ratio in electronics (or 1/SNR). It's intuitive, in that you can visualize how the random fluctuations relate to the range of the variable of interest, and it indicates approximately the expected percent error of any given value of the variable. The trouble with the standard statistical measures like correlations, percent variance accounted for, and so on is that they tend to distort the situation in favor of making it look better than it is. You'd think that a correlation of 0.8 would be 80 percent as good as a correlation of 1.00, wouldn't you? But any physicist who got an 80% correlation from his measurements of physical quantities (outside quantum physics) would immediately start dismantling his apparatus to see what's wrong with it. He'd probably do that with a correlation of 0.9, too, and even 0.99. You don't start getting good data until you're in the 0.998s (roughly 5 percent error of prediction, if I've figured that right).

Most of the difficulty in finding a control-system model for a given behavior is not so much in finding the right form of the model or the right definition of the controlled variable, but in setting the parameters of the model correctly. If you don't have the controlled variable pretty well defined, you're not going to get much of a match with the model at all. When you get the right definition, you're going to get high correlations right away with just about any old parameters -- 0.8, 0.9. Usually you start with essentially perfect control as the assumption. Then you back off on the parameters to get the best match with the actual behavior, which will usually bring you up toward 0.95 or 0.98. Then, if you still want more accuracy, you start playing with the details of the model -- whether there's slowing in the perceptual function, how much integral, proportional, or derivative component goes in the output function. That will get you into the .99s or 0.999s. If you still feel it's worthwhile to push for more accuracy, you can start in on nonlinearities, thresholds, and transport lags.

This should be the case for almost any continuous type of control. "Tracking" isn't a particularly special case; it represents control of all sorts of things, as I show in

Demo 1 where you can control half a dozen continuous or near-continuous variables including size, shape, pitch, and even a digital number that changes on the screen.

My view is that to build a new science of psychology, the first thing we have to nail down with simple experiments is the EXISTENCE of control of all kinds of perceptual variables. For this you don't need to refine the model into the stratospheric correlations. By the time you reach 0.999, the chances of this NOT being a control phenomenon have risen into the billions to one. I'd settle for a million to one. But not less. Why leave room for argument?

I think that finding exactly the correct control model can't really be done without neural circuit-tracing. There are just too many ways to accomplish the perception and control of the same variable in exactly the same way. There's also the question of what you gain by having a neurally exact model. There's no reason to believe that the next person you investigate in such depth will control exactly the same variable, or if it's the same variable, do it with the same circuitry. Up to a point there will be close similarities in the way different people do, say, pursuit tracking tasks. Beyond that point you can't generalize any more to say this is how "people" do pursuit tracking. There will be different numbers of muscle fibers involved, different numbers and placements of sensors, different mechanical advantages, different lengths of neural pathways, different visual acuities, and different settings of parameters generating different trajectories of movement, stabilities, linearities, and limits of control. The only reason to go past some reasonable level of predictive accuracy would be to characterize an individual down to the last synapse, something I wouldn't ever need to do even if I could.

I don't think that we will ever be able to say which control model is the absolutely correct one -- if only because different people will accomplish the same control tasks using different control organizations. It's more important to establish control as a phenomenon at many levels by experimental demonstration, in the manner of Galileo, and to establish methods for refining models as far as we please when such is the point. It took about 300 years after Galileo for physics to claim that further advances would be in the sixth decimal place. We aren't even 50 years into this revolution.

>I move and second that we propose Bill Powers as a speaker for
>MIT's Revolving Seminar.

I can't think of a single reason for going through such a painful and futile experience.

Best to all, Bill P.

Date: Mon Oct 05, 1992 7:31 am PST
Subject: second attempt

A MESSAGE ON THE DISCUSSION OF INFLUENCE AND CONTROL. PERSON NOT ON NET
SO ANSWER DIRECTLY. CHUCKFrom: BAKANICV@ASHLEY.COFC.EDU
To: n050024@univsc.bitnet
Subject: sorting out PCT

Dear Chuck,

It has been a long long time since I was exposed to Power's ideas. But I have enjoyed the dialogue you forwarded to me. My first reaction to the dialogue was that everyone seemed to be writing from the perspective of A (i.e., the manipulator). Discussion of B centered around whether B actively participated in the manipulation as B's attempt to control B's own environment (albeit indirectly) implying that B is also an A. That lead me to conclude that the A-B positions were a false dichotomy and everyone is an A desperately trying to match refence signal and control their own environment. But, I rejected that notion intuitively. I couldn't quite figure out what bothered me about that scenario. I've been thinking about it for several days and I believe its the powerfulness of the explanation that both attracts me and doesn't feel right. PCT is a very self empowering theory. It locates control in the behavior of each individual. Although the choices an individual makes may be constrained by environment and the actions of others which impinge upon the environment, individuals still power to adjust behavior, reset reference signals or simply remove themselves from crisis producing situations. But, seeing the world as an infinte arrangement of interlocking closed loop control systems which produce a constance of variance to which individual organisms respond, seems at odds with the sense of control depicted in the A-B scenarios.

That lead me to wonder if we weren't all B. That idea appeals to me because I have experienced social interactions as a women (i.e., a less powerful position, subject to the attempts of others to control me). I asked myself how would my interpretation of PCT theory be different if I assume the object of interaction was not to manipulate and control, but to protect oneself from manipulation and restore prior conditions. Does it make any difference? I am still pondering this. The perspective of A has the illusion of power: manipulator, setter of reference signals, adjustor of one's own behavior. But the complexity of behavior and the constant variance of environment constructs a very different perception. Has your group discussed this dilemma? Perhaps my uneasiness is a result of stepping into the middle of an ongoing dialogue. You may already have answers to my uneasiness. Von

Date: Mon Oct 05, 1992 7:50 am PST
Subject: Greg's summary

[From Bill Powers (921005.0730)]

Greg Williams (920928) --

I don't believe that the following is quite correct:

>4. The path of learning/reorganization is a function of (possibly
>randomly generated) successive sets of changes in reference signals
>and/or input/output functions, each successive set of changes being
>made to the result of the previous set of changes, and another set
>of changes being made only if certain criteria are not met for
>ceasing learning/reorganization.

If reorganization changed a reference signal directly, as opposed to changing the organization of an output function that generates it, it would be altering a SIGNAL that is being emitted by a higher output function (except at the highest current level). But this could mean only injecting a signal into the same output path already occupied by the existing signal. It seems to me, therefore, that injection of arbitrary reference signals by reorganization would be limited to the highest level of existing control systems. To inject such a signal at any lower level would simply be to disturb existing control systems which are ALSO supplying reference signals at that lower level, with the result that the disturbed systems would alter their own outputs and cancel the effect of the reorganization-produced change.

Changing a signal does not create a permanent change of organization; it only affects content, not form, and does so only a long as the signal remains changed. True reorganization, as I have thought of it, alters parameters, not signals. As a result, of course, signals take on new relationships to each other.

Hebbian learning makes parameters depend on long-term effects of the signals passing through the network that is being reorganized; that's the Hebbian version of reorganization. For this to work in general, it must be true that there is some "best" way of handling signals. This assumption shows up in the postulate that the output of a neuron somehow "strengthens" the effects of input signals that exist at the same time as the output signal. The implication is that it is best for the organism that all signals contributing to a given neural output have the maximum possible effect. This sort of rule, I believe, is an attempt (of which I approve) to get away from a "teacher" that already knows how a neural function should be organized. But I don't think it will actually work. Any system in which there is a preferred kind of input or output function loses the ability to adapt to environments in which some other kind of input or output function is required for successful control. Hebbian learning will not work if what is needed is an input function with particular positive and negative weights on its inputs. I don't think it is generally true that an increase of effect of an input is always better than a decrease or no change.

The reorganizing effects I imagine act only on parameters, not on signals.

>Here it is assumed that memories of environmental disturbances
>count as environmental disturbances, so that, for example, the
>disturbance of being called a "pig" would, as a memory, continue to
>act as a disturbance perhaps for a long time after the sound waves
>had dissipated, with similar results as if "pig" were being

>repeated over and over again by the disturber.

"Memory" is a somewhat ambiguous term, because it's often used to mean any effect that makes a signal persist beyond the termination of the input that produced it. I think you're using it that way here. If a perceptual function has a long decay time, the perceptual signal will continue to exist after the lower level signals at the input of the perceptual function have disappeared. The perceptual signal indicating the occurrance of "pig" is not itself, as I'm sure you realize, a word, but only an indication THAT a particular word has occurred. That's the only sense of "memory" that fits your proposal above. Memory that requires recording and associative retrieval would not constitute this sort of memory; in order for it to serve as a disturbance, the same memory location would have to be addressed again and again, and the system involving it would have to be operating in the imagination mode. The retrieved signal is then under control, and is no longer equivalent to an arbitrary external disturbance.

>6. At any time, the criteria for ceasing learning/reorganization >are functions of reference signals and input/output functions at >that time.

I still have a problem with both aspects of this. In effect, you're saying that intrinsic or critical reference levels can be set by (a) a signal in the system to be reorganized, and (b) the FORM of a function. I think it's incumbent on the proposer to show how a system that worked like this would be organized.

I have a feeling that to achieve the effect you're thinking of, you need both a built-in reorganizing system of the type I propose and some other type of reorganizing system that can accomplish changes in parameters based on this-lifetime experience. It would help if you could give examples of the second kind of reorganizing, so we could judge where it does something that a learned hierarchy can't do.

I have located a place to get the two photoreproductions done at $11 per page. I'll take the materials in today.

Best,   Bill P.

Date: Mon Oct 05, 1992 7:52 am PST
Subject: Competing behavior

[From Bill Powers (921005.0830)]   Martin Taylor (921002.2230) --

>I never intended to imply that the competing behaviours had any
>connection internal to the hierarchy other than that they all had
>references to which a higher-level ECS contributed. They certainly
>should not inhibit one another internal to the hierarchy,

If the competing behaviors do not inhibit each other inside the hierarchy, this means they must be mechanically interfering with each other outside of it. But as only one of a mutually-exclusive set of behaviors can be occurring at one time, I fail to see how there can be more than one such behavior going on. If I take the car rather than the motorcycle, how can taking the motorcycle even occur, outside the hierarchy? How can there be any other behavior to compete with the actual behavior of driving the car? As far as I can see, the hierarchy can accomplish its goal either by driving the car or by using the motorcycle; whichever one is going on is the ONLY one. The other behavior simply does not occur.

I'm not convinced that "competing behaviors" means anything. Perhaps you can spell out the mechanism you mean in more detail.

>I see you still don't buy the notion that not all controllable
>percepts are at any one moment controlled, and that most ECS gains
>either are very low or the ECS is in some way disconnected from the
>physical world.

I'm not sure what I said that gave you this idea.

Best,   Bill P.

Date: Mon Oct 05, 1992 7:52 am PST
Subject: Returned mail: User unknown (fwd)

Sorry it took me so long to reply. I was hard at work completing the CT text glossary. It was interesting to hear about your study. I am curious about your methodology and how you think it is analogical to PCT. >From what I could gather, it seems to be a fairly close fit. Your question about the difference of PCT from other theories has started me thinking about ways to demonstrate it. Of course, that just leads to other aspects to bring in, like the benefits resulting from PCT, and how they might be applied. I would find your views of how PCT does fit your methodology very interesting. Professor Robertson at NEIU did a study on grade control involving students in the CT course. It called for students to set goals for grades on chapter tests and track achievement. You might find that study interesting to what you are doing. As for me, I am finishing up my 120 for graduation this term. More specifically, I am "poised" for alterations that may be necessary in my life experience portfolio under the BOG Program, and I need math and science credits in CLEP. The other major goal this term is to make up my mind about graduate work toward a PhD or work on a psychological management program (which would incorporate PCT). I have extensive management experience in a civil rights career from which I took an early retirement in 1989 when I began a new one in psych. University life is very attractive which may make graduate school more interesting - particularly if work on my program could be part of it. Practicality seems imperative to whatever I am doing; working with adults to achieve optimal success is my main interest. I've been living my program and hope to use my reaching "optimal success" (hopefully, I will do that) as my model for offering it to others. I do need to earn a few dollars to supplement my annuity pretty soon, and, I am hoping that will be incidental to graduation and "making up my mind".

By the way, where are you? I am not proficient enough at this system yet to tell from your address. ?"?

Date: Mon Oct 05, 1992 7:57 am PST
Subject: Giving in

From Greg Williams (921005)

>Bill Powers (921004.0800)

>I guess what I object to in your proposals that we can "guide"
>reorganization in other people is the assumption that the guider can
>arbitrarily set this as a goal and then just carry it out. The whole
>discussion of "influence," from your end, seems to be setting up the
>influencer as an independent agent, with the influencee simply
>reacting as the influencer wants. This puts all the volition and
>agency into the influencer and ignores that of the influencee.

I have taken pains to try to show, on a PCT basis, that one person CANNOT "arbitrarily" control JUST ANY of his/her perceptions depending on another person's acting a certain way.

>It
>really seems to me to be an attempt to get around the interactive
>nature of control and go back to the old picture of the organism's
>behavior as a passive consequence of external events.

Sorry about your confusion on this. I do think your confusion is mainly due to your own active processes, and has little to due with what I actually say. Nevertheless, I'm willing to point out our areas of agreement once again.

>When we influence the actions of others in systematic ways, we succeed
>(when we do) only BECAUSE they are independent controlling agents, not
>DESPITE that fact.

I agree.

>You can't make a person reorganize without inducing critical error into that

>person, or creating a situation that does it for you.

Fine. I'd still like to hear how you think more mundane learning works.

>The immediate reaction
>of the other person will be to find SOME way to counteract this most invasive
>act, or to negate or escape from the situation that produces the error, by the
>most direct and immediate means possible including setting the situation back
>to the way it was. The other person isn't going to just wait around passively
>while you, to suit your own goals, start causing pain or suffering. You can
>set problems all you like, but there's no constraint on the other person to
>solve them according to the rules you had in mind, or to seek the particular
>solution you consider the "right" one, or even to be interested in the problem
>in the first place. When you use this method, you're in constant danger.

No doubt.

>I objected that you can predict the outcome of reorganization only
>when you arrange the environment so that only one action or learned
>method of control will in fact correct the critical error. You took
>this to mean that there IS a way of guiding reorganization -- just do
>what I described. But I offered this as an objection because it is so
>unlikely that you could ever do this with a human adult, a peer,
>without some great source of power behind you.

Remember that I was only claiming that the four types of control involving what others do are POSSIBLE according to PCT, with an explicit caveat that any of the types of control might be DIFFICULT in particular situations. I am sympathetic to your notion that guiding reorganization isn't easy. How about guiding more mundane learning, though?

>I claim that in order to establish the conditions under which this
>kind of influence (and many other kinds) could be even apparently
>effective, you would have had to establish a background of power over
>the other person and over the environment, power denied to the other
>person. In many of your examples of means of influence, I believe you
>are assuming (perhaps without knowing it) surrounding conditions that
>amount to establishing a far greater degree of coercive control over
>the other person than the degree of influence indicated by the simple
>example. This is particularly true when you speak of influencing or
>guiding another person's reorganizations.

I agree with you that a background of "power" relations can (sometimes, at least in the short-run) aid the success of control depending on what others do. But so can a background of "exchange" relations. Now, maybe you want to count ANY exchange relation as a species of coercive "power" relation. That obscures the symmetric-asymmetric distinction I was trying to make awhile back, but it would certainly be in keeping with your thoroughgoing "it's ALWAYS me vs. the world" philosophical view. (That reminds me of a story. Between my junior and senior years in college, I worked in DC for consumer rights advocate Ralph Nader. Toward the end of the summer, I began to get fed up with his acting as if -- I would phrase it now -- HIS reference signals were the best ones for EVERYBODY. So I demanded an audience, and asked him, "Do you really think that all the people who think they're NOT getting screwed by the corporations are wrong?" He simply said, "Yes." For Ralph, "freedom's just another word for nothin' left to lose.") At any rate, treating "exchange" relations -- including those where all parties claim they have benefitted from the exchange -- as "power" relations, I AM assuming such background relations. But I don't think, as you seem to, that all such background relations are in some sense heinous.

>But it is true even of such
>innocuous means of influence as rubber-banding. How are you going to
>keep the other from seeing what you're doing? Or, after seeing it,
>from objecting to it or deliberately changing the goal simply to
>frustrate your attempt at influence?

As I pointed out way back, you can only model what the other person wants, and then test that model and revise it if necessary. As we've agreed, control depending on what another person does CANNOT be arbitrary. You would be well- advised to NOT rubber-band with someone who tells you he/she doesn't want to do it.

>You propose as one effective means of controlling another:

>>3. A arranges B's environment so as to trigger learning
>>/reorganization in B's control system resulting in actions which A
>>wants to perceive.

>The natural reaction of B is to prevent A from arranging B's
>environment in that way, perhaps by sinking a knife into A or by
>driving a car past A's house and letting loose with an Uzi. This
>method of control dooms itself to failure. That's why I don't treat it
>as a serious proposal: it can't work, except temporarily.

IN SOME CASES, that would be "the natural reaction" of B. But not in cases exemplified by B going to A -- and paying A money -- to "teach" B to swim. Of course, I agree that there is a (heinous?) background lurking covertly: A won't "teach" for free, and the pool is behind a security desk.

>There is another reason I don't treat it seriously: that is because
>reorganization will NOT result in the actions which A wants to
>perceive, but in whatever actions B decides may work best for B. You
>assume that A can arrange the environment so that only ONE action can
>succeed.

No, a finite SET of actions. Skinner called such a set an "operant."

>But you overlook the action of changing the environment back
>to its original state, you assume that it's POSSIBLE for A to make
>this arrangement without any interference from B, and you assume that
>A can anticipate all possible outcomes of reorganization (such as
>suicide). I think these assumptions are unwarranted, in general.

I don't overlook, assume, assume, or think the above. As I pointed out before, A needn't know beforehand ANY solution to the problem he/she sets for B. A only needs to be able to recognize a solution when (and if) B "gets" it. Explicitly: control via "guiding" learning/reorganization is NOT GUARANTEED TO WORK EVERY TIME. I know that ever so well, based on homeschooling our kids. I also know ever so well, on the same basis, that control via "guiding" learning/reorganization is DOES NOT FAIL EVERY TIME.

In the absence of compelling evidence backing up your "it's basically impossible" arguments, I must again agree to disagree with you.

I suggest that you interact with Gary Cziko, rather than me, for at least a couple of days regarding "outside" influences on learning/reorganization. It appears to me that you are more reasonable in your claims when replying to Gary than when replying to me on the very same points, perhaps simply because you don't want to surrender a millimeter to me. Sometimes, higher-level goals can get in the way of understanding, as you previously pointed out to me regarding ideologically driven arguments.

Best wishes,   Greg

P.S. Did you get the corrected pages for the arm paper? Are they OK? How did photographing the figures work out? Have you sent the paper to SCIENCE yet?

Date: Mon Oct 05, 1992 8:40 am PST
Subject: A tracking task: help needed

We are setting up a student project for which we need some help. Any suggestions on how to tackle the problem are gratefully received.

There is a video tape showing a moving person. On his/her arm, there is a light source, so that what we need to do is to track the light spot (a classical tracking task). The output of the control system is simply the xy coordinates of the light spot as a function of time (which will be further processed, later).

A computer simulation of this is perfectly acceptable, although what we really want is to use the real data, that is, the video tape.

Any suggestions? Thanks in advance,

Marcos Rodrigues   mar@uk.ac.aber

Date: Mon Oct 05, 1992 9:47 am PST
Subject: reorganization

[From Bill Powers (921005.0900)]   Greg Williams (921003) --

>Note that I have included BOTH "learning" AND "reorganization" in
>3. Do you think we DON'T need to postulate something "less" than
>reorganization in cases which Gary has raised, such as learning how
>to multiply (where it appears that (1) the learner's control system
>is altered and (2) there is no critical error triggering the
>learning)?

Why does it appear that there is no critical error "triggering" the learning of multiplication? I've looked at this sort of learning the other way around: multiplication requires learning to control a new pattern of perceptions; therefore reorganization must occur. If reorganization occurs, what sort of critical variable departs from its reference level and is restored by successful learning? This kind of critical error would have to occur not specifically for learning multiplication, but for learning anything of that sort, under the conditions that exist at the time of learning.

There are some obvious answers and some not-so-obvious answers. The obvious answers involve the coercive atmosphere of school, in which not learning means confinement, disapproval, and punishment. For most children the motivation for learning to multiply consists of all the things that are done to them until they do learn -- more specifically, the emotional consequences and the underlying state of internal disorder resulting from the things that are done to them at school and at home until the time when they demonstrate to someone else's satisfaction that they know how to multiply.

The not-so-obvious answers would concern mainly those few children who learn to multiply because they think this is a totally neat new skill and they can't wait to find out how to do it. Critical variables, as I have maintained from the start, are not confined to the vegetative functions, although they are easiest to understand in terms of biochemical functions that underlie hunger, thirst, and so on. All that is required to define a critical variable is that it represent a condition of the organism that is genetically specified, that it can be detected by a built-in critical-variable detector, and that it lead to reorganization when it departs from its inherited reference level.

The learning of multiplication would certainly not itself be a critical variable; that is, there is no built-in reference level for learning multiplication or any other specific cognitive skill. But there could be built-in detectors and reference signals that are satisfied when ANY new skill is acquired, when ANY previously- experienced perception is brought under skilful control, at ANY level of organization. That sort of critical reference level could be inherited, so it could be operational from early during gestation. It would account for a generalized urge to learn new skills, at any level including the highest one currently under construction. And it would satisfy my demand that evolution not be required to function on the basis of future conditions of the environment that have only ephemeral existence, such as a particular cultures, languages, or most-admired sets of skills.

I've proposed one such critical variable, which is simply the absolute magnitude of the error signal in any control system. It doesn't matter what this error signal is about; all that matters is that it become as small as possible.

If teaching relied on the built-in reference signals for critical variables having to do with learning new skills, it would not be necessary to force learning by induction of critical errors of irrelevant kinds. Simply not knowing how to reproduce an observed skill would be enough, if the skill seemed interesting. Unfortunately, the educational system is so fundamentally coercive that few students have time to discover the innate joy of learning. They are too busy learning how to avoid the penalties of NOT learning that the educational system arbitrarily imposes.

Best,   Bill P.

Date: Mon Oct 05, 1992 10:21 am PST
Subject: Re: Learning and memory

[Martin Taylor 921005 11:30]   (Bill Powers 921005.0730)

This is in response to the discussion between Bill and Greg.

To start with, both have produced a set of principles, but neither gave a functional description of an ECS and its place in the hierarchy. To have a clear discussion of learning, we have to have agreement on what the hierarchy looks like, at minimum. So here is what I have understood to be more or less agreed. If any of this is wrong, it affects the discussion and should be put right.

As I understand the hierarchy, the unit of which all is built is the ECS. The ECS has two kinds of input function:

(i1) a perceptual input function, which combines all the many sensory inputs according to some algorithm and produces a scalar value called the perceptual signal. The word "sensory" includes the possibilities of direct input from sensor systems and the perceptual signals produced by the perceptual input functions of other (lower) ECSs.

(i2) a reference input function, which combines action output signals from other (higher) ECSs. Typically, this function is considered to be a simple summation, but it could be any algorithm. The output range of the reference input function must not exceed that of the perceptual input function. The output of the reference input function is a scalar value called the reference signal of the ECS.

The perceptual signal is compared with the reference signal to produce a scalar value called the error signal. The error signal is transformed by an output function, typically but not necessarily an integrating amplifier, to produce an output signal.

The ECS has two kinds of output that are distributed to other ECSs (except that in the case of the lowest level ECSs the output goes directly to muscles of other effector systems. In each case there is the possibility that the output is weighted differently in transmission to each individual destination: (o1) A set of perceptual outputs that consist of the perceptual signal possibly mulitplied by some weight. These are the sensory inputs of other (higher) ECSs. (o2) A set of action outputs that consist of the output signal possibly multiplied by some weight. These are the reference inputs of other (lower) ECSs.

The ECS has the possibility of storing one or more values of its perceptual signal and of using a stored value in place of its reference signal (episodic memory).

The ECS has the possibility of using an action output signal as one of its sensory input signals (imagination).

The hierarchy consists of ECSs linked only by the connection of the (possibly weighted) perceptual signals of lower ECSs to sensory inputs of higher ECSs, and of the (possibly weighted) action output signals of higher ECSs to the reference inputs of lower ECSs. Each kind of link is one-to-many.
===================

If the foregoing is a correct description of an ECS and its place in the hierarchy, what opportunities are there for learning? The following all seem plausible:

Within a set of ECSs already in existence:

(1) Alteration of the perceptual function.
(2) Alteration of the perceptual-sensory link structure.
(3) Alteration of the output function.
(4) Alteration of the action-reference link structure.
(5) Alteration of the reference input function.
(6) Alteration of the content of the internal memory of the ECS.

And
(7) Incorporation of a new ECS into the hierarchy.

There are subclasses:
(1a) (3a) and (5a) Modification of the parameter values of the function
(1b) (3b) and (5b) Alteration of the form of the function.

(2a) and (4a) Alteration of the connection weights
(2b) and (4b) Alteration of which ECSs are linked.

All in all, this makes 12 logical possibilities for types of ways the hierarchy can learn. Not all are effective, and the discussion has generally focussed on only a couple.

Hebbian learning is a term usually denoting topologically smooth changes of parameter values in a combining function. There are many forms, but generally speaking they all involve changes based on some measure of goodness of the present set of parameter values in relation to the data input to the function. One way is to make more extreme a pattern that creates a large output, and to make less extreme a pattern that produces a small output. Another way is to move the input parameters in such a way as to increase the output when some "teacher" asserts that the input pattern is one to which the function "should" give a large response.

Within the ECS hierarchy, Hebbian learning has usually been discussed as it applies to the perceptual input functions (1a), and then usually when the input functions are seen as nonlinear compressions of weighted sums of their inputs. If the hierarchy consisted only of ECSs having this form of perceptual input function, the perceptual side would be a classical multilayer perceptron, and Hebbian learning would suffice to allow it to perceive (and hence possibly to control) any describable partitioning of the sensory input data space. Other forms of perceptual input function, possibly involving delayed sensory inputs, are possible and may be necessary for the control of dynamical percepts such as sequences.

Hebbian learning could apply in classes 1a, 3a, or 5a. Because the success of control does not depend much on loop gain if the loop gain is high enough, one would not expect it to be very important in 3a, except in adjusting parameters such as integration time constants. It might be important in 5a, which affects the relative strength of different higher-level ECSs on the reference signal of the ECS in question. No discussion (that I remember) has considered this possibility, or how it might work if it happened at all.

Hebbian learning could apply in 2a and 4a as well, but numerically the results would be indistinguishable from 1a and 5a applied in a different ECS. The reason these possibilities are listed is that there is a potential question about the location of responsibility for the alterations of weights. "What does the ECS know, and when does it know it?"

Hebbian learning cannot be relevant to structural alterations in the hierarchy (1b-5b), because the learning is not topologically smooth. This is the province of "reorganization." It also cannot apply to classes 6 or 7, which involve discrete events that change either the content or the structure of the hierarchy.

(Powers)
>Hebbian learning makes parameters depend on long-term effects of the
>signals passing through the network that is being reorganized; that's
>the Hebbian version of reorganization. For this to work in general, it
>must be true that there is some "best" way of handling signals. This
>assumption shows up in the postulate that the output of a neuron
>somehow "strengthens" the effects of input signals that exist at the
>same time as the output signal. The implication is that it is best for
>the organism that all signals contributing to a given neural output
>have the maximum possible effect. This sort of rule, I believe, is an
>attempt (of which I approve) to get away from a "teacher" that already
>knows how a neural function should be organized.

This is indeed one form of Hebbian learning, but the motivation is not to get away from a teacher, so much as to provide the maximum discrimination among input signals that is consistent with the variation in stimulus patterns. Usually, weights on low-valued inputs decrease when weights on high-valued ones increase, so as to sharpen the discrimination.

> Any system in which there is a preferred kind of

>input or output function loses the ability to adapt to environments in
>which some other kind of input or output function is required for
>successful control.

True. And I think that your "levels" of ECS accommodates that. But Hebbian learning can operate within different kinds of input function, provided that the function is such that small changes in parameter values cause small changes in the function's behaviour.

I had intended to carry on this posting with a discussion of reorganization possibilities, but I think it is long enough, and I have spent too long on it, already. Maybe later. But if you accept my categorization of learning possibilities, there seem to be 5 kinds of reorganization, of which perhaps 1 or 2 may be useful in learning. And then there are episodic memory learning and the introduction of new ECSs. So there are many possibilities for ways in which the hierarchy might change to develop new skills.

When there are so many possibilities, it makes good sense to see what happens when only the most prominent are used. In this case, I think that the most likely candidates for effective learning are (in no particular order), 1a, 2+4b (together), 6, and 7. These can be verbalized as "What can I perceive" "What can I do to control what I perceive" "What have I perceived" and "I can do nothing right--let's try something completely different."

Martin

Date: Mon Oct 05, 1992 12:06 pm PST
Subject: Eye-head, cheap insects and sub-symbolic

From: Oded Maler 921004 3 quick notes:

1. In Nature (around 24.9.92) there was an article about a model that predicts head-eye movements during driving.

2. There was a recent announcement in comp.robotics concerning very cheap insect robots such that even independent researchers can afford to play with (less than 10$ each). If there's interest I can post. Their perceptual system is very primitive, but hackers can extend it.

3. I cam across some work of Smolensky where he claims to unify sub-symbolic and symbolic notions in a common framework. Does Bill P. have any ideas/opinions concerning this work (which is done in the neighborhood) ?

--Oded
Oded Maler, LGI-IMAG (Campus), B.P. 53x, 38041 Grenoble, France
Phone: 76635846 Fax: 76446675 e-mail: maler@vercors.imag.fr

Date: Mon Oct 05, 1992 12:06 pm PST
Subject: STELLA Intro Package

[from Gary Cziko 921005.0313 GMT]

A few days ago a received an intro package for STELLA II from High Performance Systems.

It is a very nice package indeed. The highlight is two diskettes with the full STELLA II program, except that you cannot save files or print. Included are about six models from fields like biology, math, physics, chemistry and economics which you can run and modify and run again as you wish. The mini-manual provides a tutorial for building your own simple model involving animal population growth.

For anyone with a Macintosh who wants an introduction to dynamic modelling, I highly recommend this package, especially since it's free (curiously, the price of the full program is not mentioned anywhere that I can find it, although I know that students on this campus can get it all for $65 if they are enrolled in certain courses that use STELLA).

The representative that corresponded with me is Steve Peterson. He can be reached in Hanover, NH using e-mail <x0858@applelink.apple.com>, fax (603.643.9502), or phone (603.643.9636).

I have yet to get information on TUTSIM to compare it to STELLA. Greg Williams provided only a snail mail address and I have yet to get around to printing my letter. If Greg has a fax or e-mail address for TUTSIM, I'd be much obliged.--Gary

P.S. Perhaps Bill can let us know what he found out about the possibilty of modifying STELLA so that real-time data can be used.

Gary A.

Date: Mon Oct 05, 1992 12:10 pm PST
Subject: PLEASE SEND ME YOUR ABSTRACT FOR A PAPER - QUICKLY

Please send title and abstract of a paper by October 7, 1992 for:

A PERCEPTUAL CONTROL THEORY OF MACRO STRUCTURES

to

Charles W. Tucker
Department of Sociology
University of South Carolina
Columbia SC 29208
V- (803) 777-3123
F- (803) 777-5251
H- (803) 254-0136
N050024 @ UNIVSCVM

Other important events at this meeting are:

A session on Kuhn' Self Theory organized by Bob Stewart with presentations by Carl Couch, Tom McPartland, Clark McPhail and Chuck Tucker

A Presidental address by Clark McPhail

MIDWEST SOCIOLOGICAL SOCIETY MEETINGS
APRIL 7-10, 1993
HYATT REGENCY
CHICAGO, ILLINOIS

Date: Mon Oct 05, 1992 12:26 pm PST
Subject: Re: Competing behavior

[Martin Taylor 921005 14:00] (Bill Powers 921005.0830)

>If the competing behaviors do not inhibit each other inside the
>hierarchy, this means they must be mechanically interfering with each
>other outside of it.

Yes, that's right. If they don't inhibit one another in the world, and are provided with reference signals by the same higher-level ECS, they will act simultaneously.

>But as only one of a mutually-exclusive set of
>behaviors can be occurring at one time, I fail to see how there can be
>more than one such behavior going on. If I take the car rather than
>the motorcycle, how can taking the motorcycle even occur, outside the
>hierarchy?

On other occasions, the car may not have been available. On this occasion, whichever behaviour had the lowest impedance would inhibit the other(s). It took reorganization for the hierarchy to develop in such a way that taking the cycle was a possible way of reducing the perceptual error, equally with taking the car.

This discussion revolves around whether it is a property of an ECS that it "knows" what it is doing when it provides an output. I have been taking the position that each ECS is blind, in that all it knows is that it provides an output when it has a non-zero error. If the hierarchy has been effectively reorganized, that output will reduce the error more probably than not (whether it actually does is dependent on the momentary state and disturbances of the world and on other control systems that affect components of its feedback circuit, but the ECS knows nothing of that). You seem to require the ECS to know what lower-level ECSs are doing and to choose which one to activate. I will turn back on you your own frequent question. What percept does it control to do that?

Martin

Date: Mon Oct 05, 1992 1:39 pm PST
Subject: Why 99%, hierarchical perception

[From Rick Marken (921003.1100)] Martin Taylor (921003)

>The statistics can give a clue. If the
>residual variance is proportional to x and to y, logarithms are the likely
>answer. If it is more or less independent of x+y, then mutiplication is
>probable. What then? If a logarithm can be developed in one part of the
>hierarchy, is it not likely that it can be done in another part? Then perhaps
>we should look for logarithmic relations elsewhere in the hierarchy. But
>if it looks more as if the answer is multiplication, then a host of different
>relations seem reasonable candidates as the CEVs in other situations.

Martin, trust me (I'm a Dr.). The "Las Vegas" approach to doing research can now be chucked. No need to analyze residuals, no need to look for patterns in noise. Just build the models and do what it takes to get them to behave exactly (to the level you like -- say less than 5% error) like the real system. Tom Bourbon (921002.1050) -- also a Dr.-- and Bill Powers (921004.0800) -- not a Dr. but, even better, a genius -- explain the approach extremely well. The method of PCT is modeling. The method of moribund psychology is statistics.

Bill Powers (921003.0600) --

A wonderful post on the perceptual basis of HPCT. Let me just put in a quick plug for an approach to exploring perception that I describe in my soon to be rejected "Hierarchical control of perception" paper. I have described this on the net (maybe) but I do so again in the hopes of jogging some of those other minds out there for some suggestions -- and maybe develping some new HPCT demos.

The method described in my paper is very simple -- numbers alternate back and forth on the screen so, with time going from top to bottom, what is presented is:

```
5
     7
8
     2
4
     6
etc...
```

The rate of alternation can be varied by the observer. When the rate is very fast (the max possible on the computer -- say about 15/sec -- all you can see is the numbers -- their configuration. When you slow down the rate of alteration you get to a point where you see the numbers "move" back and forth, like the "phi phenomenon". Slow it doen even more and you can start to see the "sequence" -- you can tell that 5 comes before 7, then comes 8 and then 2, etc. The sequence can by perceived only when the alternation rate is about 4/sec. If you slow it down even more you can "see" that there is a rule underlying the sequence -- if number on left >= 5 then number on right is odd, else number on right is even. The observer can know this rule in advance but cannot perceive it (at least, this here observer can't) until the alternation rate is about .25/sec.

This demo is not supposed to be earth shattering; it is just an attempt to provide a helpful way for people to examine their own perceptions. It seems to me that it might help someone understand what it means to perceive a "configuration", a "transition", a "sequence" and a "program". The variations in rate help you "isolate" the perceptions and see that it is possible to have a low level perception (like transition) that implies the possibility of a higher order perception (a sequence) and still not be able to perceive the higher level percept (until the rate is slowed). It is interesting that it seems to take longer to perceive "higher order" perceptions but this, in itself, does not imply a hierarchical relationship between the perceptions. Bill's logical test -- that you can't perceive certain things unless you can perceive their constituents -- seems like a better basis for claiming hierarchy. Interestingly, the transition perception goes away when the rate slows too much -- it obviously depends on other things too -- such as distance between the numbers. But, since we do see the sequence even though transition is gone, it seems like sequence perception does not depend on having a perception of transition -- but it does depend on having a perception of configuration (since it's a sequence of confugurations).

I think there must be "perceptual demos" of this sort that might help to demonstrate some ways to look at perception (using our best and most accessible lab -- our own brain) and see why at least some of us PCTers think the H in HPCT represents quite a bit more than an opinion.

Best regards   Rick

Date: Mon Oct 05, 1992 4:52 pm PST
Subject: Reorganization; tracking; Hebb; competing behaviors

[From Bill Powers (921005.1330)]   Greg Williams (921005) --

What I'm being so stubborn and obtuse about (I can see how it looks from your end) is really simple: I see learning as a skill of an organism. A rock does not learn, no matter how long you try to teach it or how you disturb it. The mechanisms of learning -- those that decide when it is necessary to learn, those that decide when learning has been successful, and those that actually create changes of neural connections and synaptic weightings -- belong exclusively to the organism. There is nothing the outside world can do to make these MECHANISMS work any differently. This is just like saying that the mechanisms of muscle contraction are properties of the organism that no outside environment can change: muscles contract in response to driving signals because of their biological construction, not because of the world in which such contractions take place. So it is with learning: learning is a function of a biological system with the necessary equipment.

The kernel of this "necessary equipment" is what I call the reorganizing system. This system represents a built-in skill of the organism; it works according to its inherited design, over which the outside world has no control, and on which neither the outside world nor the inside world has even any influence.

The argument we're having is failing to converge because we're talking about two different things. What I'm talking about is the fundamental organization of the organism; what you're talking about is how this organization interacts with particular external events and processes. I'm talking about how a muscle works; you're talking about what that muscle will do under various kinds of driving signals.

Because of the built-in capacity to reorganize, and the built-in criteria that go with it, this particular piece of equipment has an especially important function: it makes the organism (the human one, at least) infinitely adaptable to external circumstances. From the standpoint of the organism, IT DOES NOT MATTER what particular control processes have to be learned. The organism has no built-in preference for controlling any specific variable in any particular way -- not, that is, among those variables controlled by the learned part of the organization. The only unchangeable requirement is that the critical variables be maintained near their respective reference levels. The world outside the organism has absolutely no influence on that requirement; it is absolute. Critical reference levels are the only ones in the organism about which we CANNOT say that they are set as they are in order to accomplish something else (during one organism's lifetime).

"Mundane learning" works because of critical error. All learning does, that involves a change in any functions. Of course there are also other processes having nothing to do with reorganization that are called learning: executing systematic search patterns, memorizing facts, memorizing and imitating sequences of actions, executing program-like algorithms that employ present-time information from outside the organism, and so on. Those are either memory phenomena, or are themselves learned control processes acquired initially through reorganization.

My point is not that organisms have no influences on each others' behavior, or even that these influences can't be manipulated to have intended effects. It is that these are all surface phenomena, acquired through the control actions of a reorganizing system. Each organism, as suggested today by "von", simply controls for its own perceptions, fending off any external influences that don't produce effects it already wants. This applies to manipulators just as much as to the manipulated; a manipulator is just trying to make the world look the way the manipulator wants it.

I am also not rejecting the idea of interaction. Organisms can ask for help (in numerous ways), and get it. They will reject the very same "help" if it is given without being requested. Organisms can ask for instruction and information, and get it. They will reject the same instructions and information if it was not requested -- provided they are given any choice in the matter. The demand for learning comes from inside the organism, not from outside it. There is no such thing as "teaching" as a transitive verb. All you can do it make it possible, or necessary, to learn. It isn't the teacher who guides the learner, but vice versa.

Ah, well. We've run this subject into the ground for now.

Marcos Rodregues (921005) --

I take it that you want to track the light on the moving arm in order to obtain a record of the movements. Just how automatic a project do you want this to be? If you want a device that will optically track the light and put out signals that indicate its position in two or three dimensions, this could become quite a project.

Here's an illustration:

Let the light be a photodiode modulated at, say, 1000 Hz. To detect its position you need to form an image of it with a lens in a small telescope, focused on a cluster of 4 photocells, two for the x direction and two for the y direction. The photocell signals go through an AC amplifier and phase-sensitive detectors synchronized with the 1000 Hz signal that modulates the photodiode. This eliminates ambient light. Each pair of photocell signals enters a differential amplifier that yields a position error signal. The position error signals for x and y become reference signals for servos that move the telescope in x and y. Thus the telescope is always aimed so that the pairs of photocells are illuminated equally. The servo feedback signals are used as the indicators of position. As the arm moves, the servos make the telecope track the light. If you want depth information you have to use two telescopes. A-to-D converters will sample the servo feedback signals, and from the geometry you can calculate the actual x-y (-z) position of the light.

This can also be done more cheaply. Build a frame with three degrees of freedom that is pivoted at the shoulder and at the elbow. Strap the frame to the arm. Put linear potentiometers on each pivot. Put voltages on the pots, connect their outputs to 3 A/D converters, and let the computer calculate the arm segment configuration.

To do this off a video tape, you can put an electronic gate on the position of the light. The gate would be divided into four regions, and the video signal within each region would act like the signals from the photocells. The gate would be moved (in time-delays) instead of moving the telescope; otherwise the principle is the same. But again, not a simple project.

Gary Cziko (921005.0313 GMT) --

>P.S. Perhaps Bill can let us know what he found out about the
>possibilty of modifying STELLA so that real-time data can be used.

So far I haven't heard back from Hammond.

Martin Taylor (921005.0730) --

That was a nice summing-up of the basic HPCT architecture. I think that all of the candidates for reorganization are valid. Some, such as the linkages from one level to another, may be most actively reorganized during development (I think we learn new perceptual organizations at a decreasing rate later in life). Also large-scale changes in linkages of all kinds probably decrease in frequency (during development, neurotaxis seems very important; this is a more drastic way of making new connections than merely altering the weights of synaptic connections).

My main difficulty with Hebbian learning is that it seems unmotivated; that is, it's so local that I don't see how it could be related to critical variables. The only test of it that I know about, in terms of modeling, was done by a graduate student in Harry Klopf's lab at Wright-Patterson Air Force base; to keep all the weights from going to maximum or zero, all sorts of kludges had to be introduced. The basic model by itself didn't work. Maybe others have made it work better since then.

>This is indeed one form of Hebbian learning, but the motivation is
>not to get away from a teacher, so much as to provide the maximum
>discrimination among input signals that is consistent with the
>variation in stimulus patterns. Usually, weights on low-valued
>inputs decrease when weights on high-valued ones increase, so as to
>sharpen the discrimination.

So this is a "normalized" model? That would partly take care of the problem of making some weights decrease. I think a workable version of Hebbian learning would have to be pretty far advanced over the original notion of just "increasing the strength" of connections.

I still don't understand the mechanism by which competing behaviors prevent each other from happening outside the organism. Just saying that a behavior has a "low impedance" doesn't explain much. You still leave some of the systems with unsatisfied errors inside the organism, meaning very large outputs. Why aren't those large outputs producing large efforts?

>I have been taking the position that each ECS is blind, in that all
>it knows is that it provides an output when it has a non-zero
>error.

So have I.

>If the hierarchy has been effectively reorganized, that
>output will reduce the error more probably than not (whether it
>actually does is dependent on the momentary state and disturbances
>of the world and on other control systems that affect components of
>its feedback circuit, but the ECS knows nothing of that).

I don't see what "probably" has to do with it. The only way for an error to be small is for the output to be in the state that brings the perceptual signal nearly into a match with the reference signal and keeps it there. If there's any uncertainty about whether the error has been reduced, you're going to have a very low-gain control system, incapable of opposing disturbances to any interesting degree.

>You seem to require the ECS to know what lower-level ECSs are doing
>and to choose which one to activate. I will turn back on you your
>own frequent question. What percept does it control to do that?

The higher one knows a perception derived from copies of the lower perceptions (plus, in general, some uncontrolled perceptions). A system that acts by selecting which lower level action to perform will cease to look for an effective choice as soon as the first one is found that corrects its error. This would be at least a category control system, where selection might be in terms of the name of the class of perceptions. Should I take the "bicycle" (quick check for a perception in that category) or (failing that) a "car?" (something in that category is found and the reference level is satisfied). This implies a priority sequence. Another way would be simply to look for something in the class "transportation." Then the search wouldn't be prioritized. If you happen to spot the roller skates first, that ends the search. When any suitable perception is found, the error is zero and the scan ceases. It is then up to other systems to do something with the "bicycle" thing or the "car" thing that has been spotted: drive it away, wash it, sell it, paint it.

Best   Bill P.

Date: Tue Oct 06, 1992 6:22 am PST
Subject: percept

[From: Bruce Nevin (Tue 92106 08:35:57)]

I am troubled by the noun "percept" that you have introduced, Martin, and have been using for some time. Maybe you can clear it up for me.

I know you defined it carefully at the outset, a number of months ago, but I'm reluctant to search through disk archives. (And I don't have a Mac with your Hypercard stack!) It seems to denote something outside an ECS to which the perceptual signal within the ECS corresponds. Perhaps there is a "percept" corresponding to the reference signal as well, and maybe even for the error signal?

At one level, my concern is that this term cannot legitimately be used when speaking from the point of view of the ECS, and that to do so is to mask certain kinds of possible assumptions (of reader or writer), and consequently to court confusion. (Bateson talks at length about epistemological problems that arise from errors of logical type.)

At a second level, my concern is that some unjustified reification may have crept in around the use of this term "percept," however carefully it was originally intended.

It seems to be useful to abbreviate otherwise cumbersome locutions. Perhaps it would be a good thing to spell out explicitly what locutions are intended to be abbreviated, as distinct from others for which we might be tempted to use it but ought not.

Or is it all crystal clear, and I'm just out of touch with CSG dialect?
        Bruce   bn@bbn.com

Date: Tue Oct 06, 1992 7:32 am PST
Subject: Re: percept

[Martin Taylor 921096 10:45]   (Bruce Nevin 92106 08:35:57)

>I am troubled by the noun "percept" that you have introduced, Martin,
>and have been using for some time. Maybe you can clear it up for me.

>I know you defined it carefully at the outset, a number of months ago,
>but I'm reluctant to search through disk archives. (And I don't have a
>Mac with your Hypercard stack!) It seems to denote something outside an
>ECS to which the perceptual signal within the ECS corresponds.

Interesting comment. I was not aware of having changed terminology, and I would have to go back through my postings to see exactly where I have used the term "percept". But if I were to use it today in a posting, it would mean exactly "perceptual signal" (or possibly a stored memory of a perceptual signal that could be used as a reference signal). I might use it to refer to a desired perceptual signal, which would be one that matched a reference signal, but I doubt that I have used it for the reference signal itself. I would not use it to refer to the error signal. And I do not have a concept of a percept as being anything outside the ECS.

This question raises a subtle issue that this question raise, and one that has supported an ongoing theme for some time. The question is more or less "What is it in the outer world that corresponds to the perceptual signal in the ECS?" I have called that relationship among outer world variables the CEV (Complex Environmental Variable). Sometimes I may have spoken of the percept of the CEV or something of the kind. There was an exchange with Bill a little while ago on what he labelled "Three Stages of Satori": (1) I control the CEV, (2) I control the percept, (3) I control the CEV, realizing it is really the percept.

The CEV can be described to and possibly observed by a third party, whereas the percept cannot. The great problem of The Test comes from the difficulty of accurately determining what is the CEV for the subject's percept. That is the issue that is ongoing with Rick Marken about the statistics of his area-perimeter study, and by extension the value of statistical analysis in psychological experiments generally.

The CEVs, being in a mirror world, are at once personal and public. They form the link between people, but they are determined by private perceptual functions that may not be overtly known to the perceiver, and hence cannot be accurately described to another person. But the variables and relationships that are the stuff of the CEV exist in the world (unless you believe in solipsism), and if I push something that forms part of one of your CEVs, it alters the corresponding percept. So in that sense, language sometimes may lead me to talk as if the percept was in the world, like the CEV that mirrors it. But as soon as you want to be precise and technical, the percept (a present or possible perceptual signal) is in the ECS and nowhere else.

I know that's not all clear, but the whole mirror-world thing is only gradually becoming less vague in my own mind.

Martin

Date: Tue Oct 06, 1992 9:12 am PST
Subject: Percepts; statistics

[From Bill Powers (921006.0945)] Bruce Nevin (921006.0835) --

>I am troubled by the noun "percept" that you have introduced,
>Martin, and have been using for some time. Maybe you can clear it
>up for me.

I agree with Martin's definition. In my 1960 paper with Clark and McFarland (A general feedback theory of human behavior, Part I; Perceptual and Motor Skills _11_,71-88) I defined it this way:

"A _percept_ is the basic unit of experience. It is that "bit" of perception which is self-evident to us, like the intensity of a light, or the taste of salt.

"A _variable_ is always a combination of two classes of percept. One class contains percepts which _do not vary_; by these percepts we keep track of the "identity" of the variable. The other class contains percepts which _do_ change; these percepts carry information about the "magnitude" of the variable. "Magnitude" is used here in its most general sense, including the meanings of "intensity," "size", or any other word for the general class of variable attributes." (p. 3-4 in LCS).

All this fancy definition-making did NOT lead to Principia Gubernatoriae, although I had, then, ambitions in that direction. But it's clear that Clark, McFarland, and I started right out defining the world from the standpoint of the perceiver, with the percept being the basis for everything. What I had in mind, I think, was to define control theory without any reference at all to an objective outside world. I obviously gave up on that as being impractical. As Martin Taylor (921006.1045) puts it so neatly:

>(1) I control the CEV, (2) I control the percept, (3) I control the
>CEV, realizing it is really the percept.

Rick Marken (921003.1100) --

I think that Martin has been saying lately that statistics is used improperly in psychology (predicting individual behavior from mass measures, etc.), but that it still has uses in understanding signa-to- noise ratio, observational error, and the like. I think this dead horse will stay dead now, at least locally.

I take a somewhat different tack, which is that the world of experience doesn't seem statistically uncertain. Whatever the actual noise level of neural signals (and it isn't zero), it is low enough on the time scale of perception and action that it plays no important part in HPCT. The world looks smooth, continuous, and sharp down to the limit of resolution -- which is all we care about. This tells us that most neural signals involved in ordinary behavior have magnitudes large enough to make noise unimportant.

Uncertainty that crops up in behavior is, I think, to be attributed to conflict, not noise. Only when we're poised on the knife-edge between conflicting wishes, or when unusual circumstances require us to see in the dark, smell something at the threshold of detection, or hear something on the bring of inaudibility, does the underlying noise level make any difference. What people call "choices" I would call "conflicts." When there is no conflict, you simply act to produce and maintain the world you want, choice being unnecessary.

Best to all, Bill P.

Date: Tue Oct 06, 1992 10:02 am PST
Subject: re: percept

[From: Bruce Nevin (Tue 92106 12:22:23)] (Martin Taylor 921096 10:45) --

You probably haven't changed your use of "percept" since introducing it. Evidently, the confusion is mine. I can see how it came about. Whenever you said "percept" I assumed that you meant something different from "perceptual signal": neither expression is any more or less unwieldy than the other, so the obvious motivation that I would have in your place would be to make a terminological distinction.

Because of your background and training, I assumed "percept" must be a neologism in the jargon of psychology, now being applied by you to PCT. Because of its shape, it looks like a product nominalization to me, denoting the product or result of the process ("perceiving") named by the verb, e.g. on the following analogy:

conceive : concept :: perceive : percept

It was because of this confusion that I took "percept" to refer to something other than the perceptual signal itself (neural current in a nerve fiber in an ECS). I could make sense of the word only by supposing that you were reifying a referent for the perceptual signal--yet I couldn't believe that of you.

If there is a chance that others might be similarly confused, mightn't it be better to stick with "perceptual signal," "reference signal (stored memory of a perceptual signal)," and "desired" or "potential perceptual signal (matching a reference signal)," respectively, for the three senses you've identified?

Bruce bn@bbn.com

Date: Tue Oct 06, 1992 11:31 am PST
Subject: percept

[From: Bruce Nevin (Tue 92106 13:21:19)]

(Bill Powers (921006.0945)) --

So, clearly, Martin did not introduce the term. But in your usage it seems to evoke our subjective experience ("qualia") rather than a neural current in a nerve fiber:

>A _percept_ is the basic unit of experience. It is that "bit" of
>perception which is self-evident to us, like the intensity of a light,
>or the taste of salt.

In any case, I hope you can see that the typical reader, whether dipping into LCS or into CGS traffic, is unlikely to bear the technical definition in mind and might well be misled by the conceive : concept analogy. I am not the most hasty or careless reader you will encounter, and I forgot. It perhaps calls for frequent "by the way" reiteration of its definition, like acronyms, when we really must use it instead of "perceptual signal" etc.

>What I had in mind, I think, was to define >control theory without any reference at all to an objective outside >world. I obviously gave up on that as being impractical.

The effort founders I think on the fact that you want to communicate. In order to communicate, we assume prior agreement about what appear to be CEVs in the world, and in the process of communicating we verify and negotiate these and other agreements.

It seems important to us to understand those CEVs as being in the public world, external to all communicating parties (including the CEVs that constitute the communicating parties ourselves). This craving seems to arise because such agreements are a prerequisite for cooperative action, and a great many of our controlled perceptions are influenced by actions of others (cooperative or not).

But public stability is an assumption, a meta-agreement that is part and parcel of the "prior assumption" just mentioned. But this is to say only that we maintain reference signals for having agreements about CEVs in the public domain, and when discrepancies result in error signals we act to re-establish those agreements. This seems to me to be what compels negotiation and motivates the search for agreements. For the sake of the perception of public stability, these agreements once attained are attributed to the world as being knowledge of the world.

Agreements, and the cooperation that they enable, may even be sometimes more important to us than the "fidelity" of public-domain CEVs with personal perceptions. We may tolerate discrepancies. Sometimes people may compartmentalize or encapsulate systems of perceptions so as to avoid perceiving inconsistencies (Rokeach: _The Open and Closed Mind_).

Oh dear, I wasn't going to say much. I'll never catch up at this rate!

Bruce bn@bbn.com

Date: Tue Oct 06, 1992 1:48 pm PST
Subject: Re: truth vs cooperative agreement

[Martin Taylor 921006 16:30] (Bruce Nevin Tue 92106 13:21:19)

>Oh dear, I wasn't going to say much. I'll never catch up at this rate!

Wow! I'm really glad you said as much as you did. I think you made a foundational statement in the last part of your posting. Just for luck, I want to play it back for you. Why? I like it so much I want to see it again. It says so much in so few words.

Sorry, troops and bandwidth conservers...
===============

>Agreements, and the cooperation that they enable, may even be sometimes
>more important to us than the "fidelity" of public-domain CEVs with
>personal perceptions. We may tolerate discrepancies. Sometimes people
>may compartmentalize or encapsulate systems of perceptions so as to
>avoid perceiving inconsistencies (Rokeach: _The Open and Closed Mind_).

Martin

Date: Tue Oct 06, 1992 1:55 pm PST
Subject: Re: Reorganization; Hebb; competing behaviors

[Martin Taylor 921006 11:30] (Bill Powers 921005 13:30)

One place where I, and I suspect Greg, have not come to agreement with you is in the limitation of the initiation of reorganization to error in intrinsic variables--at least this is how I interpret your term "critical error." You have, once or twice, allowed that a form of critical error is some kind of integration of error over all or part of the hierarchy, but for the most part I read you as identifying critical error with errors in variables I would characterize as body-integrity variables. My view of the mechanism of reorganization is (I hope) the same as yours, but the triggering, for me, is a sustained and particularly a growing error in any ECS. What is reorganized is the output links from that ECS, either by sign changes or by delinking or forming new links. That (local) reorganization may well trigger further reorganization in either lower or higher ECSs by affecting their ability to control, but apart from this avalanche effect, I don't subscribe to a concept of global reorganization.

In the Little Baby project, our first experiment will be with local reorganization, with a Poisson rate that is a monotonic function of the instantaneous value of (e*e + k*e*de)/G where e is the current error, de the difference in error from the previous sampling moment (a surrogate for the derivative), k an arbitrary parameter, and G the gain. Chris has implemented this, but we haven't tried it yet. The reason we put G in there was that you reported bad effects if you switched the sign of a link in a high-gain ECS.

(Incidentally, have you progressed any further in your experiments on reorganization mechanisms?)
=========================

>My main difficulty with Hebbian learning is that it seems unmotivated;
>that is, it's so local that I don't see how it could be related to
>critical variables.

Let's take that in two parts. The motivation for Hebbian learning is that it allows for gradient learning. Typically, the network doing the learning is intended for classification of some more or less complex input data space, such as identifying handwritten characters according to their letters of the alphabet. Less often, but importantly, it is intended to reduce the apparent degrees of freedom of incoming data by performing what amounts to a non-linear principal components analysis. In both cases, the network develops its weights in response to the statistics of its input, where input may be only the incoming data or may include "training" (teacher's identification of "this" pattern as an instance of Q, for example). As an example, we might have a multilayer perceptron with an input data space of, say, 20 by 30 pixels each of which may have a vlaue between zero and unity, representing the amount of "ink" in the pixel in an image of a handwritten character. There may be few or many input units, all of which see these 600 inputs. The input units connect to few or many hidden units, and the hidden units connect to 26 output units. The teacher decides beforehand which letter is to be output by each of the 26 output units, and the idea is that if Q is presented, the Q output unit should have a unity output and the others should have a zero output.

The network is provided with many examples of handwritten characters imaged in this 20 x 30 space, and for each, the teacher specifies what the outputs of the 26 output units should be. A training algorithm (of which there are many) alters the weights of all the units according to the errors of what the output units should have perceived and what they did perceive. With luck, the next time that character is presented, the set of 26 outputs will be nearer to the desired pattern of one unity value and 25 zeros.

In a teacherless system, the changes in weights are such as to exaggerate differences among input patters, leading units within a level to decorrelate their outputs, or to converge to a common output, depending on their initial sets of weights, the statistics of the input patterns, and the local interconnections among units. Each additional level increases the range of the complexity of discriminations that the net can make. One layer can discriminate only linearly discriminable patterns (I'm talking about units that simply produce a non-linearly compressed weighted sum of their inputs, not more complex units such as radial basis functions). Two layers can discriminate arbitrary connected regions, and three can discriminate arbitrary finite sets of regions (if I remember correctly). This is true whether the network is trained by a teacher or not.

My idea about the network of ECSs is that the connected perceptual input functions form precisely a multilayer perceptron if each is of this simplest (non-linear summation) type. Unlike a normal multilayer perceptron, the ECS has a local criterion for whether it is contributing to the network as a whole--its error. In a normal neural network that is only an S-R one-way system, responsibility for any output of the net is distributed over a large number of internal units. In a control net, each ECS knows whether it is controlling its percept effectively, and responsibility can therefore be localized. If this is so, then the error can be a criterion for Hebbian alteration of the perceptual input weights, and, we might hope, for robust and rapid learning when it is compared with the performance of a normal perceptron learning algorithm.

When an ECS fails to control, it may be because its output are ineffectively connected, because there is conflict with another ECS, or because the CEV it is perceiving is inherently uncontrollable given the effectors that link the hierarchy to the world (like the rising of the sun). The first case demands reorganization (and possibly Hebbian changes in the output weights), the third demands Hebbian changes in the perceptual input function (and possibly reorganization), and the second case seems balanced between the two. But it is not clear whether anything internal to the ECS can be used to discriminate among the three cases.

In the Little Baby project, the second (probably) experiment will be to start with a random set of output connections, and use Hebbian learning on the perceptual input functions to see whether the baby can learn to perceive the world in a way compatible with its (fixed randomly assigned) outputs. The third, and most interesting experiment, if the first two succeed, is to see whether reorganization and Hebbian learning can be used together, actions adapting to perceptual input functions that change toward more controllable forms. That's a little way off yet, but I could imagine that at least some results from one or two experiments might provide a Christmas present for the group. I hope so.

============= >I still don't understand the mechanism by which competing behaviors >prevent each other from happening outside the organism. Just saying >that a behavior has a "low impedance" doesn't explain much. You still >leave some of the systems with unsatisfied errors inside the organism, >meaning very large outputs. Why aren't those large outputs producing >large efforts?

When one of the mutually competing behaviours acts so as to provide a percept that satisfies the higher-level ECS's reference signal (taking the car gets me closer to my destination), the reference signals for the other would-be competitors are reduced. As we approach the destination, the reference signal for perceiving myself to be taking the bike is reduced, as is the error in that I perceive myself not to be taking the bike. I suspect that the imagination loop is active in such cases, too, in that when I take the car, I can imagine myself thereby arriving at the destination, immediately eliminating the error signal in the "bike-taking" ECS. But suppose the car will not start. Then the higher-level reference is not being satisfied, and as you say, the output of the "taking bike" system increases, to the point where it overtakes the "taking car" system and inhibits it. I get out of the malfunctioning car, kick it for luck, and get the bike out. (That's if I don't also have a high-gain ECS with a reference for perceiving the car to work properly, regardless of whether I am using it at the moment. If I do, I may call the repair truck, which inhibits the higher system that requires me to reach my destination.)
================

>>If the hierarchy has been effectively reorganized, that
>>output will reduce the error more probably than not (whether it
>>actually does is dependent on the momentary state and disturbances
>>of the world and on other control systems that affect components of
>>its feedback circuit, but the ECS knows nothing of that).
>
>I don't see what "probably" has to do with it.

The world is unstable, with unpredictable disturbances. Sometimes pushing on something that always worked one way now makes it move the other way. There are disturbances and conflicts that can cause perverse effects even if the world actually is working the way you expect.

This was more a defensive paragraph, against those who would correctly point out that no control system always finds its percepts moving in the way that they usually do, given a particular output signal. If they did, you might as well have a simple pre-planning system of the type you call "cognitive". Control systems exists as a defence against the unpredictability of the world.

>The only way for an
>error to be small is for the output to be in the state that brings the
>perceptual signal nearly into a match with the reference signal and
>keeps it there.

I disagree. There are lots of situations in which error is small by chance. I return to the degrees of freedom argument. Almost all controllable percepts are not, at any moment, being controlled. But when their error exceeds tolerable limits, control may be shifted to them. This implies that most of the time, most percepts have tolerable errors as a consequence of other behaviour or the luck of the world.

>If there's any uncertainty about whether the error has
>been reduced, you're going to have a very low-gain control system,
>incapable of opposing disturbances to any interesting degree.

Again, I disagree, and this is the basis of the statistics argument that I am sure will flare up many tims before we come to agreement on it. You tend to take the long-term view of what happens as a control system comes to a stable state with (as I perceive an inderlying assumption) an invariant reference signal. I tend to concentrate on the transient behaviour of the feedback loop, as I feel that the world is perpetually disturbing percepts at all levels, thus changing reference signals at all levels, sometimes abruptly. I don't think that many ECSs get much opportunity to come to a steady state, though they may be close a lot of the time. You can have high-gain systems that sometimes misjudge the data on which they base their error signals, so long as the misjudgment doesn't last too long or happen too often.

Anyway, as I said, "probably" was for accuracy rather than as an assertion that the probability was far from unity.

Martin

Date: Tue Oct 06, 1992 3:59 pm PST
Subject: Percepts; reorganization; competing behaviors

[From Bill Powers (921006.1530)] Bruce Nevin (921006.1321) --

>But in your usage [of "percept"] it seems to evoke our subjective
>experience ("qualia") rather than a neural current in a nerve
>fiber:

The model is intended to assert an equivalence: that all objects of awareness are neural currents in nerve fibers. These neural currents, in turn, are representations, or in Wayne Hershberger's more suggestive term, "realizations", or in drier modelling terms, "functions" of an underlying order or Boss Reality knowable to us only in the form of neural currents. Or, of course, neural currents derived through more complex functions from the neural currents of lower order. This includes all we are aware of, including our thoughts about our experiences.

>The effort founders I think on the fact that you want to
>communicate. In order to communicate, we assume prior agreement
>about what appear to be CEVs in the world, and in the process of
>communicating we verify and negotiate these and other agreements.

Yes. The only time it is profitable to keep the "pure" model in the forefront is in the privacy of one's own investigations. Then it's an essential aid in avoiding taking things for granted -- by exempting, for example, the analysis one is currently developing from the status of being "only neural currents." Communication between people who have become easy with the basic assumptions of the model, however, takes on a different flavor from ordinary communication that assumes a common objective world. While it may seem that such people are talking about controlling physical things and relationships in an external physical world, the meanings being evoked in the participants include a knowledge that the behaving system is really acting on its own perceived world, and that the person describing the situation is shaping this description of the behaving system with the understanding that it is the describer's perceptions that are really being indicated. All descriptions of this sort are simultaneously theoretical propositions being tested for agreement with other people's perceptions of the situation, with observations always carrying the understood preface, "Here is how it looks to me, with what evidence I have to justify my description:"

>But public stability is an assumption, a meta-agreement that is part and
>parcel of the "prior assumption" just mentioned. But this is to say only
>that we maintain reference signals for having agreements about CEVs in the
>public domain, and when discrepancies result in error signals we act to
>re-establish those agreements. This seems to me to be what compels
>negotiation and motivates the search for agreements. For the sake of
>the perception of public stability, these agreements once attained are
>attributed to the world as being knowledge of the world.

Again, yes. When we're communicating, the only sane thing to do is to assume that our understanding of the agreement is the same as others' understanding. But the CT-aware person never accepts agreement at face value. As Martin Taylor might point out, and in other contexts has done, there are far more degrees of freedom in the world than we have under control, even during communication. What I agree to is a monster object with more degrees of freedom than my means of expression have, and I never know the extent to which your agreement refers to exactly those degrees of freedom I am considering.

Some really nice statements in this post. Glad to see you getting up to speed again.

Martin Taylor (921006.1130) --

>One place where I, and I suspect Greg, have not come to agreement
>with you is in the limitation of the initiation of reorganization
>to error in intrinsic variables--at least this is how I interpret
>your term "critical error." You have, once or twice, allowed that
>a form of critical error is some kind of integration of error over
>all or part of the hierarchy, but for the most part I read you as
>identifying critical error with errors in variables I would
>characterize as body-integrity variables. My view of the mechanism
>of reorganization is (I hope) the same as yours, but the
>triggering, for me, is a sustained and particularly a growing error
>in any ECS.

I don't want to put any a priori limits on what can amount to a critical variable. I've proposed myself that error signals in the hierarchy can amount to critical variables, in that minimizing them is a goal that is compatible with the requirement that reorganization not be dependent on the nature of the current environment (or acquired knowledge of it). I'm sure that as we come to understand more about how perception gets organized, and how levels of control develop, we will find other such aspects of the hierarchy that are context- independent and can qualify as critical variables of the inheritable sort. Such things would certainly make the building of the hierarchy more efficient than my simple overall error-driven random reorganization.

I concur with your intuition about "a sustained and particularly a growing error" as driving reorganization. My experiments with reorganization, which continue sporadically, have convinced me that there must be a strong rate-of-change component in the computation of critical error signals. Just by trial and error, I found that the most reliable method (so far) entails (roughly) making the rate of reorganization depend on the rate of change of absolute critical error TIMES the absolute critical error, so that the changes get smaller as the remaining error diminishes. I realized only afterward that this amounts to taking the derivative of the SQUARE of the error, which neatly takes care of assuring positive values and that the result will give the best least-squares fit to the optimum solution. You might try this in addition to your (e*e + k*e*de)/G. In other words, try d(e^2)/G.

Another tip. Rather than just make random corrections based on critical error, I have found that it's best to randomly change a parameter that determines the direction and speed of changes in parameters. Between reorganizations, the changes then continue on each iteration. This is strictly analogous to the E. coli method of locomotion, in which, between tumbles, the organism continues to swim in a straight line, continuously altering its relationship to a radial gradient.

If there is an array of parameters d[i], I define an array delta[i]. On each iteration, d[i] has delta[i] added to it. This is like moving in a straight line in hyperspace. As long as this movement continues to reduce the measure of critical error by a sufficient amount, reorganization is suppressed. Sooner or later, however, the movement of the parameters in hyperspace will pass the point of closest approach to producing zero critical error, and critical error will begin to increase. Then reorganizations are commenced, which alter the entries in the array delta[i] between positive and negative limits, at random. Between reorganizations, the critical error is monitored, so you can tell whether a reorganization left the error getting larger or getting smaller. If it makes the error start getting smaller, you've found a good direction in hyperspace (not necessarily the best, but good is good enough). I haven't yet played with varying the number of samples used to determine whether the error is getting smaller -- right now I accept any decrease as reason enough to stop reorganizing, and any increase on a single iteration as reason enough to reorganize. Refinements are obviously possible.

I have tried normalizing delta[i] so the sum of squares is 1. This makes delta[i] into a unit vector. I think this works a little better than just using raw deltas, but it may not be worth the computing time, all things considered.

I still see some problems in switching signs, but I could be doing something wrong. It shouldn't make as much difference as it does.

>>My main difficulty with Hebbian learning is that it seems
>>unmotivated; that is, it's so local that I don't see how it could
>>be related to critical variables.

>Let's take that in two parts. The motivation for Hebbian learning
>is that it allows for gradient learning.

That's not exactly what I meant by "unmotivated." You actually supplied the "motive" I meant when you said

>The teacher decides beforehand which letter is to be output by each
>of the 26 output units, and the idea is that if Q is presented, the
>Q output unit should have a unity output and the others should have
>a zero output.

Here, the teacher is supplying the missing motive for reorganization. The teacher already knows that there is a "Q" present, and reorganizes the network until it, too, reports a "Q". The critical variable is the output of the network as perceived by the teacher. The critical reference level is "Q". The critical error is in the teacher, who is acting as a reorganizing system. The teacher's output acts to cause reorganization in the network, which is terminated only when the teacher experiences zero critical error. The things being reorganized in the network have nothing to do in themselves with "Q"-ness -- they're just weights. The network itself doesn't care which output it produces in the end. If the teacher wanted it to indicate "A" every time there is a "Q" present, it could be made to do so (just relabel the output lines).

So this is just like my model for reorganization, except that the teacher is not inside the system but extraneous to it.

Your second example is about "teacherless" training:

>In a teacherless system, the changes in weights are such as to
>exaggerate differences among input patterns, leading units within a
>level to decorrelate their outputs, or to converge to a common
>output, depending on their initial sets of weights, the statistics
>of the input patterns, and the local interconnections among units.

If there really are no criteria of error involved here, then such a system will just converge to a state that expresses its properties, which are fixed. This would not fit my idea of reorganization; it's more like a fixed algorithm. I would guess that it's less capable of learning than the kind with a teacher. But maybe this is an important kind of network anyway, in that it would have many possible output states and could provide ABITRARY discriminations, for acceptance or rejection by a reorganizing system. I suspect, however, that the meanings of the discriminations are subject to a lot of interpretation by the human beings who are looking at the results. As you've described this sort of system, I can't see anything that would constrain it to making USEFUL discriminations. Aimless complexity isn't necessarily useful.

I think that in the overall picture, a teacher is necessary. But the criteria this teacher should use must be relevant to the system getting reorganized, not placed externally in a different organism. My reorganizing system is a teacher that uses criteria relating to the functioning of the system itself: the teacher wants all the control systems together to have as little total error as possible.

>In the Little Baby project, the second (probably) experiment will >be
to start with a random set of output connections, and use >Hebbian
learning on the perceptual input functions to see whether >the baby
can learn to perceive the world in a way compatible with >its (fixed
randomly assigned) outputs.

I can tell you already that this will work, using my method of reorganizing outlined above, with up to 20 independent control systems controlling through 20 shared environmental variables. The critical variable is total squared error across all systems. Reorganization is applied globally, not based on each system's error.

I didn't have the patience to let 50 systems converge, but it looked as if they were headed that way (that's 2500 weights being reorganized). There's no reason it shouldn't work with any number of systems. When the output connections have fixed signs (although chosen at random), and you just reorganizing the input weights (n weights per system when there are n systems controlling n environmental variables), convergence always occurs and with reasonable efficiency. And this is the worst case, because there are no spare degrees of freedom. Of course it's possible that the choices of output weights will preclude a solution (I think). I haven't run into that yet, although some choices make convergence definitely slower.

>The third, and most interesting experiment, if the first two
>succeed, is to see whether reorganization and Hebbian learning can
>be used together, actions adapting to perceptual input functions
>that change toward more controllable forms. That's a little way
>off yet, but I could imagine that at least some results from one or
>two experiments might provide a Christmas present for the group.

That would be a nice present.

RE: competing behaviors

>When one of the mutually competing behaviours acts so as to provide
>a percept that satisfies the higher-level ECS's reference signal
>(taking the car gets me closer to my destination), the reference
>signals for the other would-be competitors are reduced.

What reduces them? I still don't see how you can be riding a bicycle and driving a car at the same time.

>Then the higher-level reference is not being satisfied, and as you
>say, the output of the "taking bike" system increases, to the point
>where it overtakes the "taking car" system and inhibits it.

Where are these outputs acting, outside the system? What kind of outputs are they? Come on, Martin, you're waving your arms.

>The world is unstable, with unpredictable disturbances. Sometimes
>pushing on something that always worked one way now makes it move
>the other way. There are disturbances and conflicts that can cause
>perverse effects even if the world actually is working the way you
>expect.
>
>This was more a defensive paragraph, against those who would
>correctly point out that no control system always finds its
>percepts moving in the way that they usually do, given a particular
>output signal. If they did, you might as well have a simple pre-
>planning system of the type you call "cognitive". Control systems
>exists as a defence against the unpredictability of the world.

The levels exist to eliminate the need for reorganization or random processes. If control of certain things often entails switches of sign, a higher level system monitoring the relationship between direction of action and direction of effect will be acquired to make the required switch without any trial and error, immediately (where "immediately" means, apparently, in about 0.4 sec, according to Rick's experiments). Ordinary variations in disturbance or parameters do not require any adaptation from a control system; it just acts as it usually acts.

>>The only way for an error to be small is for the output to be in
>>the state that brings the perceptual signal nearly into a match
>>with the reference signal and keeps it there.

>I disagree. There are lots of situations in which error is small

>by chance.

Nope. If external forces bring the error to zero by chance, the output of the system will drop to zero. If it doesn't, it will CREATE an error. The output is always in the state that brings the perceptual signal to a match with the reference signal and keeps it there, even when that required state of the output happens to be zero. "Chance" has no effect on error signals. Error signals are affected by physical variables: outputs and disturbances.

>Anyway, as I said, "probably" was for accuracy rather than as an
>assertion that the probability was far from unity.

That makes quite a difference.

Best to all,   Bill P.

Date: Wed Oct 07, 1992 3:37 am PST
Subject: Scratching the surface

From Greg Williams (921007)

Gary: FAX number for TUTSIM Products: 415-325-4801

>Bill Powers (921005.1330)

>What I'm being so stubborn and obtuse about (I can see how it looks
>from your end) is really simple: I see learning as a skill of an
>organism. A rock does not learn, no matter how long you try to teach
>it or how you disturb it. The mechanisms of learning -- those that
>decide when it is necessary to learn, those that decide when learning
>has been successful, and those that actually create changes of neural
>connections and synaptic weightings -- belong exclusively to the
>organism. There is nothing the outside world can do to make these
>MECHANISMS work any differently.

This is getting supremely delicate. I agree with you that the mechanisms of an organism's learning are solely the organism's, but I also contend that significant particulars of the way those mechanisms work depend on BOTH the history of the organism's innards (and, by implication, the histories of its ancestors' innards) AND the history of its interactions with its environment (ditto for ancestors).

>This is just like saying that the mechanisms of muscle contraction are
>properties of the organism that no outside environment can change: muscles
>contract in response to driving signals because of their biological
>construction, not because of the world in which such contractions take place.

I agree that the mechanisms of neural transmission remain unaffected by an organism's environment (barring physiological trauma), just as the muscle contraction chemistry does. But just as muscles can hypertrophy when "driven by" exercise, so can control structure organization become altered when "driven by" environmental problem-setting (especially by other controllers). A particular manifestation of hypertrophy/alteration depends significantly on BOTH innards AND outards.

>So it is with learning: learning is a function of a biological system with
>the necessary equipment.

It is a function of that AND disturbances due to an independent environment.

>The argument we're having is failing to converge because we're talking
>about two different things. What I'm talking about is the fundamental
>organization of the organism; what you're talking about is how this

>organization interacts with particular external events and processes.
>I'm talking about how a muscle works; you're talking about what that
>muscle will do under various kinds of driving signals.

The reason why I'm talking about the PARTICULARS of learning is because those are what I believe many investigators and laypersons think are important to explain and deal with. But my point that, in general, the "fundamental organization" of an organism in the here-and-now does NOT depend solely on the innards. I simply don't see how it could, going all the way back to a consideration of the basic math (the diff. eqs.) in the PCT model. Disturbances are inputs; inputs affect the PCT (with reorganization or not) variables' trajectories.

>Because of the built-in capacity to reorganize, and the built-in
>criteria that go with it, this particular piece of equipment has an
>especially important function: it makes the organism (the human one,
>at least) infinitely adaptable to external circumstances. From the
>standpoint of the organism, IT DOES NOT MATTER what particular control
>processes have to be learned. The organism has no built-in preference
>for controlling any specific variable in any particular way -- not,
>that is, among those variables controlled by the learned part of the
>organization. The only unchangeable requirement is that the critical
>variables be maintained near their respective reference levels.

This is a fine argument for the idea that the true "guiding" of reorganization to particular states which satisfy the INTERNAL criteria (after all, the end states of reorganizations appear quite NONrandom!) is the work of the environment. I suspect that there are internal processes at work in the guiding, too, but that is beside this point.

>The
>world outside the organism has absolutely no influence on that
>requirement; it is absolute. Critical reference levels are the only
>ones in the organism about which we CANNOT say that they are set as
>they are in order to accomplish something else (during one organism's
>lifetime).

Yet I suspect that it is possible to override the inherited c.r.l.'s. "Death before dishonor," "my country over my life," etc. At least the SETTINGS of c.r.l.'s or the loop gains for their associated circuits, I think we are agreed, appear to be modifiable within one's lifetime (and, I would add, in ways influenced by the environment). I.e., "wanting to be accepted by peers" comes to outweigh "wanting to survive," and the platoon members start running up to the enemy machine-gun nest.

>"Mundane learning" works because of critical error. All learning does,
>that involves a change in any functions.

This is quite a remarkable hypothesis, at least to my ears, after hearing for so long that reorganization is a gut-wrenching experience. It seems to me that you must be careful not to set yourself up for criticism similar to that applied to the rampant instinctism early in this century, when there were dozens of rather specific instincts postulated. Are you implying that there is an inherited critical variable corresponding to, say, wanting to do things well, which is in charge when I want to do things well? If so, then I suppose there is another c.v. corresponding to NOT wanting to do things well in charge when I don't? THIS SORT OF APPROACH EXPLAINS TOO MUCH. The succinct question I ask you about your critical variables of various sorts being involved in all sorts of "mundane" learning is: How would you test this hypothesis?

>My point is not that organisms have no influences on each others'
>behavior, or even that these influences can't be manipulated to have
>intended effects. It is that these are all surface phenomena, acquired
>through the control actions of a reorganizing system.

I think that these "surface phenomena" are quite important to many folks, and that a PCT-understanding of them would be welcomed by many. They certainly ARE "acquired through the control actions of a reorganizing system." And those control actions depend on the ENTIRE CONTROL SYSTEM INVOLVED, which loops through BOTH organism and environment.

>I am also not rejecting the idea of interaction. Organisms can ask for
>help (in numerous ways), and get it. They will reject the very same

>"help" if it is given without being requested. Organisms can ask for
>instruction and information, and get it. They will reject the same
>instructions and information if it was not requested -- provided they
>are given any choice in the matter.

So close, and yet so far. We basically agree, but then you go overboard with the "me against the world" view. If proffered aid is thought useful by a person, the person won't reject it. Rejecting it or not has nothing to do with whether the person requested it. Guessing well about what a person would be pleased to know, but currently doesn't know that they would be pleased, is the province of good teachers and counselors.

>The demand for learning comes from inside the organism, not from outside it.

THIS demand, NOW, comes from inside. How did THIS demand come to be? It arose from processes BOTH inside AND outside.

>Ah, well. We've run this subject into the ground for now.

To the contrary, I think we've only scratched the surface.

Bill, a check for $11 for photography goes out today. Now I anxiously await the rejection slip!

Best,   Greg

Date: Wed Oct 07, 1992 7:21 am PST
Subject: The behavior of perception

[From Rick Marken (921007.0800)]

I'm very busy here at work so I can just barely keep up with the mail. But I just wanted to give a quick news bulletin -- which I might be able to follow up on at lunch time. The "Hierarchical behavior of perception" paper was rejected by "Theory & Psychology". The two reviews were the nicest I have ever received. Basically, both said it was a good paper but "we already know that" (that behavior is controlled perception). I love it! So why are they still acting like they don't, I wonder?

Details at eleven -- really.

Regards   Rick

Date: Wed Oct 07, 1992 8:04 am PST
Subject: Smart levers

[From Bill Powers (921007.0800)]   Greg Williams (921007) --

>This is getting supremely delicate.

Yes, and I'm having a heck of a time saying what I mean. When I vary the words to bring out a meaning in one dimension I get into trouble with another. Part of the trouble is in trying to handle living-system (purposive) environmental effects in the same breath with non-living effects that have no goals.

Maybe the main problem is with the ambiguous phrase "how the system works." Here's an analogy that I junked last time because I was getting too long-winded.

You have a lever. By moving one end down or up you can make the other end go up or down (the fulcrum is between you and the load). Now, how does this lever "work?"

I just described that, didn't I? You can make the other end of the lever work any way you like, by using any pattern of movements of your end that will produce the intended result. The lever works exactly as you intend it to work.

But, in another sense, this has no effect at all on _how the lever works_. No matter how you move your end, the movement of the other end is still opposite to and proportional to the movement of your end relative to a horizontal plane. In lifting loads you will find a certain mechanical advantage; you have no control over that, either.

Now suppose this becomes a smart lever. It senses the downward force you're applying to your end of the lever, and it doesn't want this force to be greater than, say, one pound (a one-way control system). As long as the force is less than one pound, it does nothing. But if the force starts to rise above one pound, the lever begins to expend energy in a way that pushes the fulcrum toward the load end. If the force drops below one pound, the output relaxes and the fulcrum slides back toward its initial position. So in general, the fulcrum will be placed so the force on your end is just one pound, unless the load is light enough that less than one pound at your end will move it.

Now you find that when you lift a small load with the lever it works just as before. But if the load becomes larger, suddenly you have to move your end of the lever more to lift the load by the same distance, but you still have to exert no more than one pound of effort. You can still make the other end do anything you want -- you can still work the lever any way you like. But now you have to make larger movements to accomplish the same thing.

This smart lever is now controlling something about itself that matters to it -- the force on one end of its beam. It's doing this by changing its own properties -- not its behavior as you see it, but its properties. It's changing the mechanical advantage, from your point of view. It isn't resisting you -- you still have as much control as ever, and in fact it's making big loads easier for you to lift. But it's making sure that you don't effect a variable that it senses and wishes to keep in a certain state: less than one pound of force in a particular place.

As an intelligent controller with goals, you will naturally observe the movement of the fulcrum and figure out what's going on. AHA, a new variable to control. By varying the load that you put on the other end of the lever, while moving your end to lift it, you can, apparently, cause the fulcrum to move to any position you like. You now have control of what this smart control system is DOING -- of its ACTION. If you like to see the fulcrum in a particular position, you can find just the right load so that when you lift it, the fulcrum will slide to the position you want. Of course you must then give up control of the amount of load you're lifting -- that variable become subordinate to your goal of seeing the fulcrum in a particular place. In fact, you can't control both the amount of load you lift AND the position of the fulcrum.

Neither can you just reach out and push or pull on the fulcrum as a means of positioning it while you're lifting the load. If you do move it, it will move right back to where it was as soon as you let go (unless you move it in the direction that results in requiring even less than one pound of force at your end of the lever to lift the load). If you forget about lifting the load, you will find that suddenly you can move the fulcrum any way you want. But as soon as you go back to lifting the load, the fulcrum will push back against any effort you make to move it away from the load end of the beam. You will be in direct conflict with the smart lever.

From your viewpoint, this is a very complex behaving system; you will have to learn a lot of apparently arbitrary rules if you want to control its various aspects. From the standpoint of the smart lever, however, the situation can be summed up very simply: if there's more than one pound of force on the manipulated end of the lever, move the fulcrum toward the load until the force drops to one pound.

What this lever is doing is altering its own organization in a way that maintains a critical variable in the condition it wants. In the process of doing this, it changes its properties as an external observer experiences them through interacting with the lever. These apparent properties, however, are irrelevant to what the smart lever is concerned with; their changes are side-effects. They're important in the world of the observer relative to the observer's goals, but not in the world of the smart lever relative to its goal.

Clearly, the placement of the fulcrum by the smart lever depends on what is going on in the world outside it. The lever's properties (those that matter to an external observer) change as a function of external events. Its mechanical advantage depends on how much load there is and on whether someone or something acts to lift that load by pressing down on the other end of the lever. So you could say of the properties of the lever -- the externally visible properties -- that they are "... a function of [the lever's control system] AND disturbances due to an independent environment."

But we could also say that the lever's properties have not changed at all -- those that matter to IT. The lever is maintaining control of the one critical variable we have given it, relative to the one built- in reference level we have given it. The lever's control system continues to work according to exactly the same principles and with exactly the same effects no matter what external manipulations are carried out. The parameters of control do not change.

If you take the point of view external to the behaving system, you see many ways of influencing both the behavior and the properties of the smart lever. These, however, are all defined relative to the perceptions you are interested in, and your goals for those perceptions. You are doing this to preserve your own critical variables, where "doing" refers both to the way you have organized your own perceptions and control systems, and to the actions you carry out. In the end you manage to interact with the smart lever in a way that leaves all your critical variables in their prescribed reference states. So at that level of organization, neither you nor the lever has been disturbed in any significant way. All that has happened is a complex adjustment of behavioral variables that ends up with both sides in control, as before.

If we want to understand relationships among people, we have to attack the problem in this way. We can't get hung up on "good" influences and "bad" influences; we can't take it for granted that satisfying some people's goals (like educating children) is better than satisfying other people's goals (like the childrens' goals for what they would like to be doing). We have to see this problem of human interaction as just that, an interaction among independent systems. The laws of social interaction that we come up with must not have anything to do with PARTICULAR goals and PARTICULAR behaviors, or with side-effects of this interaction that can appear, depending on your point of view, as control of or influence on another person. From any individual's point of view, ALL interactions with others are control of their behavior or influences on their behavior. But that is true of all the others, too; that's how they see you. The people you are trying to influence or control are also trying to influence or control you, by the very actions and changes of properties you see as having been caused by you. To take any one side in an interaction is to miss the essence of what is going on: interaction.

Well, let's pull this one up the flagpole and see if anyone burns it.

>>"Mundane learning" works because of critical error. All learning
>>does, that involves a change in any functions.
>
>This is quite a remarkable hypothesis, at least to my ears, after
>hearing for so long that reorganization is a gut-wrenching experience.

This is what comes of either-or thinking. Either theres's a gut- wrenching critical error, or there's no error at all. The rate of reorganization, according to the theory I've been putting forth all this time, is a function of the amount of critical error, falling to zero when critical error becomes zero.

I wonder why so many universities have counselling for math anxiety, test anxiety, grade anxiety, and so on.

>Are you implying that there is an inherited critical variable
>corresponding to, say, wanting to do things well, which is in >charge
when I want to do things well?

If you mean a verbal cognitive system that has opinions about what is doing well and what isn't (for example, getting a good grade), of course not. If you mean that error signals themselves can be critical variables, with reorganization continuing until they are as small as possible (regardless of what they are about), then yes. The effect of error itself being a critical variable is that control systems will continue to reorganize until they are "doing well." But that is the outcome, not the goal. The goal is simply to minimize a certain class of neural signals.

>THIS SORT OF APPROACH EXPLAINS TOO MUCH.

If used carelessly, it certain does. You're going about it backward, though. This is not like attributing every motive to "instincts." The criterion is that whatever a critical variable is, it must not have anything to do with the current external world or with any learned perception. Avoiding your objection is precisely why I adopted that criterion for critical variables. It is also, incidentally, my reason for making reorganization random. Nothing inside the organism can know in advance what behavioral organization will have a harmful or beneficial effect in the current environment. Discovering that is and must be a pure trial-and-error process -- at least if you want the fewest possible ad-hoc assumptions in the theory.

Best,   Bill P.

Date: Wed Oct 07, 1992 10:23 am PST
Subject: they already know that

[From: Bruce Nevin (Wed 92107 13:19:46)]   (Rick Marken (921007.0800)) --

>The "Hierarchical behavior of perception"
>paper was rejected by "Theory & Psychology". The two reviews were the
>nicest I have ever received. Basically, both said it was a good paper
>but "we already know that" (that behavior is controlled perception).
>I love it! So why are they still acting like they don't, I wonder?

Can you point out something they do that is incompatible with "already knowing that"? A couple of f'rinstances, familiar and generally accepted methods and results whose incompatibility with CT you can succinctly demonstrate, and which are likely to fall within the specializations of your reviewers? They might notice that what they are agreeing to is not quite what you have said. Might not get the paper published, but the reasons for rejection might change in interesting ways. Incompatibility of the reasons for rejection, especially from the same reviewers, might even provide some kind of leverage with the editor.

The voice of inexperience . . .

Bruce   bn@bbn.com

Date: Wed Oct 07, 1992 10:57 am PST
Subject: blindsight article in _The Sciences_

[From: Bruce Nevin (Wed 92107 14:00:02)]

While waiting for some administrivia, just read the article "Unconscious vision: the strange phenomenon of blindsight" by Lawrence Weiskrantz (of Oxford) with Sr. Editor Karen Fitzgerald, in _The Sciences_, September/October 1992 (23-28). This suggests to me that control systems in the midbrain and hindbrain (not exclusively for visual perception) operate in parallel with those in the cortex, but without conscious awareness.

Descriptions of "guessing" responses to stimuli to which systems in the cortex (but not those in other systems) are blind are remarkably similar to descriptions of responses in tests for socalled psychic perception. This suggests to me a possible explanation for at least some of these experience, namely, that there may be perceptions controlled in the midbrain and hindbrain that are not controlled in the cortex.

In any case, the article (avowedly) bears squarely on the question of the seat(s) of consciousness in the brain, and summarizes interesting results and interesting questions, for which I recommend it.

Bruce   bn@bbn.com

Date: Wed Oct 07, 1992 12:14 pm PST
Subject: they already know that

[From Rick Marken (921007.1230)]   Brice Nevin (921007) --

>Can you point out something they do that is incompatible with "already
>knowing that"?

Well, the "experimental method" as commonly used in psychology; they don't test for controlled variables.

The paper was somewhat philosophical. I tried to make one big point explicitly -- that performance limitations (in terms of speed of behavior) may be limitations on the ability to perceive (rather than to produce the actions that produce) the results that are being controlled. I also tried to show how one can explore the perceptual (and, thus, behavioral) hierarchy. There was a lot in this paper that I thought was fairly new and interesting; alas, not the reviewers.

One would have to read the paper to get a good sense of what the reviewers missed (or got right). The interesting thing is that both reviewers said that they were "hard pressed" to find anything new in the paper; both mentioned TOTE units (missing the fact that Miller et al never understood that working versions of these "units" would control their perceptions), they also mentioned people who had presumably worked on the relation between perception and behavior (apparently ignoring the fact that the paper was about the fact that, from the perspective of the HPCT model, behavior IS perception). My efforts to make the paper more respectable by referring to work in the accepted literature (and always in a friendly way) backfired with the first reviewer who saw these references as old hat.

The second reviewer gets high marks for saying that he has followed my research for several years with great interest (I know who the reviewer is; I referred to several of his reearch findings in the paper). I just wish he had followed it with a tad more understanding (though, fair's fair, he seems to wish the same about me with respect to his research, though all I did was use his results, not his interpretations thereof).

Anyway, the reviews were friendly. They didn't beg me to rewrite and resubmit, but they were moderately encouraging. I might try to rewrite it, but it seems like a topic that might be more interesting to those who have already passed PCT 101. In other words, I'd rather leave it as is and publish it in the Journal of Living Control Systems. But we'll see; always nice to have a paper out there, waiting to be rejected.

Hasta Luego   Rick

Date: Wed Oct 07, 1992 1:47 pm PST
Subject: Give me a (smart) lever to move the psychological world?

From Greg Williams (921007 - 2)   >Bill Powers (921007.0800)

>You have a lever....

This is a lovely analogy. I believe I understand your point and emphasis (in fact, I believe I have understood all along!), and I have no major quibbles with your claims, other than that many others seem to be more interested in my "slant" on the subject than in your "slant." The big problem FROM YOUR POINT OF VIEW (as I see it) is that the analogy doesn't speak to what you have been claiming REALLY matters: CHANGES in the control apparatus. So, how did the smart lever come to be so smart? With absolutely NO environmental influence? The big problem FROM MY POINT OF VIEW comes here:

>If we want to understand relationships among people, we have to attack
>the problem in this way. We can't get hung up on "good" influences and
>"bad" influences; we can't take it for granted that satisfying some
>people's goals (like educating children) is better than satisfying
>other people's goals (like the childrens' goals for what they would
>like to be doing). We have to see this problem of human interaction as
>just that, an interaction among independent systems. The laws of
>social interaction that we come up with must not have anything to do
>with PARTICULAR goals and PARTICULAR behaviors, or with side-effects
>of this interaction that can appear, depending on your point of view,
>as control of or influence on another person. From any individual's
>point of view, ALL interactions with others are control of their
>behavior or influences on their behavior. But that is true of all the
>others, too; that's how they see you.

You think I disagree with all this. But I do not, with the sole exception that I think what you call "side-effects" can be extremely important, sooner or later, for the organism currently seeing them as unimportant.

And here:

>The people you are trying to influence or control are also trying to
>influence or control you, by the very actions and changes of properties you
>see as having been caused by you.

I DO disagree with this, if it is meant to be a general claim that everyone ALWAYS tries to control perceptions dependent on a co-interactor's actions. Certainly, many interactions DO involve simultaneous control by all parties of certain of their respective perceptions which depend on actions of the others. But, as well, there certainly exist many instances of (sometimes highly!) assymetric controlling, where party A is trying (with high loop gain) to control his/her perceptions of party B doing such-and-such, and where party B doesn't even CARE about controlling ANYTHING A is doing (and might not even be aware of A; sometimes, I must add, until it is "too late"!).

And, especially, here:

>To take any one side in an interaction is to miss the essence of what is
>going on: interaction.

I have been trying NOT to "take one side," but rather to show (based on PCT) how ANYONE who is attempting to control his/her perceptions dependent on another's actions, IN GENERAL, must be constrained IN PARTICULAR WAYS, because of the other being a controller, too.

Time to get back to issues in learning/reorganization?

Best wishes,   Greg

Date: Wed Oct 07, 1992 2:17 pm PST
Subject: Re: they already know that

(ps 921097.1400)   [From Rick Marken (921007.1230)]

>Can you point out something they do that is incompatible with "already
>knowing that"?

Well, the "experimental method" as commonly used in psychology; they don't test for controlled variables.

take an example from the psych lit.. show yr claim is true. show why they should. show what they miss if they don't. if you can't come up w/ a telling example of why experimentaly methodology is wrong, and why your way fixes (some of) the problems, that's your problem, not the reviewers'.

The paper was somewhat philosophical. I tried to make one big point explicitly -- that performance limitations (in terms of speed of behavior) may be limitations on the ability to perceive (rather than to produce the actions that produce) the results that are being controlled.

that's a reasonable point, and i'm not surprised they didn't disagree. but unless you have *some* (new) way of demonstrating it or arguing it (in their paradigm, not yours), then you haven't in fact added to the literature.

I also tried to show how one can explore the perceptual (and, thus, behavioral) hierarchy. There was a lot in this paper that I thought was fairly new and interesting; alas, not the reviewers.

but it's not interesting if one doesn't agree w/ the underlying assumptions and doesn't see how the discussion has any impact on one's own assumptions. if you showed how well-known phenomena can be analyzed differently and novelly in yr hierarchy--*that* way you can make an *argument* for your hierarchy. but that won't work if the reader must assume the hierarchy first and have no way to relate the discussion to their own ways of looking at things.

> they also mentioned people who had presumably worked on the relation between perception and behavior

why ``presumably''? plenty of people have had things to say about this.

> (apparently ignoring the fact that the paper was about the fact that, from the perspective of the HPCT model, behavior IS perception).

but see, i'm not convinced, from the kinds of things you write about, that it's coherent for you to say ``perception is behavior.'' i think i can map from your model straightforwardly into one w/ pieces labeled ``perception'' and ``behavior.'' you haven't convinced me i should consider the terms synonymous or what that would mean.

cheers. --penni

Date: Wed Oct 07, 1992 3:43 pm PST
Subject: simulating societies - conference

This likes it would be a natural for any followup work to Gatherings:

Date: Wed Oct 07, 1992 8:58 pm PST
Subject: assorted rubbage

this is ISAAC again and i have been occupied for the past week and had to delete the whole lot. next time i'll have the time to separate the wheat from the chaff.

ASSORTED RUBBAGE

1. to whoever is sending out the next closed loop, i have moved since this summer. my new address is isaac kurtzer p.o.box 4509 SFASU Nacogdoches, TX 75962 thankyou
2. to whoever can give me a SUCCINCT answer/definition, what in the johnjacobjingleheimerschmitt is a ECS ?! thankyou
3. to that person who has posted that he/she? was writing a article on skinnerian behaviorism
   a) is there a draft ready, preprint is just fine, that you could send me. i have a paper to write for tech. writing [not my spartan paper] and am basically writing a primer.
   b) could you write a short list (on the net) of books represntative of skinner's beliefs
   c) would you mind filling out a short interview for my primary source
   d) please reply, at least to tell me no thankyou
4. to C. McPhail, i figured you might be interested in knowing that i'll be using your ...maddening crowd and other works as a guide and reference for my FAT term paper on sparta/pct
5. to Gary "the player" Mobley, are you hooked on here. a no-reply i'll assume as no thankyou

Date: Thu Oct 08, 1992 3:15 am PST
Subject: entrapment

[From: Bruce Nevin (Thu 92108 06:46:08)]

This smells to me like an opportunity for the following sort of presentation:

Presenter: As we know, A.
Audience: (Yes, that's pretty obvious.)
P: And it's pretty well accepted that B.
A: (Well, of course.)
P: Research of Frooble and Whisket, among others, shows that C.
A: (Yeah, yeah, we've known about that stuff for years.)
. . .

P: From A, B, C, . . . it follows that XYZ.
A: (Suddenly waking up.) WHAT!! Let me go over that line of argument
again. Something's wrong.

Almost like the salesman who gets you saying "Yes. Yes. Yes." to a series of truisms, and then works in those affirmations that lead to your making the conclusion that you want to buy his vacuum cleaner. Except of course that your motivations are pure. And the salesman doesn't want you to wake up and reexamine the chain of affirmations that led to the novel conclusion. You, by contrast, want precisely that.

The XYZ would have to do with some consequences for method, interpretation, value of statistical analyses, whatever seems most germane that puts in the spotlight something that is glaringly wrong from a PCT perspective but invisibly business-as-usual from a conventional perspective.

Be well, Bruce

Date: Thu Oct 08, 1992 4:37 am PST
Subject: MATLAB

From Greg Williams (921008)

THE STUDENT EDITION OF MATLAB (Prentice-Hall, Englewood Cliffs, NJ, 1992; $50 for either MS-DOS or Macintosh version) is currently available at many college book stores. It is identical to Professional MATLAB (ca $700 list), a very flexible and comprehensive numerical computation program based on matrix algebra and calculus but including numerical integration of ordinary differential equations, EXCEPT that the student edition's maximum matrix size is limited to 1024 elements (32 x 32), it does NOT need a math coprocessor (but it will use one if present), and it has no fancy graphical printing facilities (but it does do both 2-d and 3-d screen plots -- the user can make a low-resolution screen dump or transfer numerical data into another program, like 1-2-3, for fancy hardcopies). Also, the student edition comes with several control-systems-design-related programs.

I think this is a great bargain for anyone who does a lot of math. Even though it doesn't allow real-time input/output and works with equations rather than block diagrams, it also might be worthwhile to consider as a common program for IBM and Mac PCTers who want to trade models back and forth. Program files for MATLAB are in pure ASCII, and work the same in IBM and Mac versions, so they could be sent via e-mail. The more experienced modelers could prepare well-commented (maybe even with ASCII "character-graphics" showing the equivalent block diagrams) program files which could be "played with" by less-experienced modelers, who could eventually gain programming expertise in a reasonably unintimidating environment. Due to the low price, students could have easy access to a library of PCT models.

Perhaps Gary will be moved to get the Mac version and see what he thinks about it. The price is right, I believe, even if it ends up being used only to tutor the kids in algebra. The hardware requirements are as follows: IBM, 320KB RAM, 3MB hard disk space; Mac, 1 MB RAM, System 6.0 or higher, 1 800KB disk drive and a hard disk or 2 800KB disk drives, Monaco 12-point font.

Just one more possibility.

Greg

P.S. No, Gary, I don't know the Prentice-Hall FAX number.

P.P.S. MATLAB versions are also available for VAX, Sun, Apollo; all use the same ASCII program files. So hardly anyone on the net need be left out of the fun if they want to join in.

Date: Thu Oct 08, 1992 7:52 am PST
Subject: Lever analogy; Rick's strategy; misc

[From Bill Powers (921008)] Greg Williams (921008) --

>The big problem FROM YOUR POINT OF VIEW (as I see it) is that the
>analogy doesn't speak to what you have been claiming REALLY
>matters: CHANGES in the control apparatus. So, how did the smart
>lever come to be so smart? With absolutely NO environmental
>influence?

In the lever analogy, the shifting of the fulcrum was meant to be analogous to a CHANGE in the system parameters resulting from reorganization (with a systematic rather than an E. coli method of "reorganizing" fulcrum position). The response of the lever at the load end to the manipulator's actions at the other end was meant to be analogous to the response of an acquired control system's output action to a disturbance of its controlled quantity (no control system is involved here in the analogy, but the relationship is that of disturbance to opposing action). The lever itself and its control system was meant to be analogous to the inherited body (the physical lever) and the reorganizing control system which has a fixed organization during a given lifetime.

>So, how did the smart lever come to be so smart? With absolutely NO
>environmental influence?

The smart lever itself is a product of evolution, which is a more basic reorganizing process with reference signals pertaining to critical variables closer to the level of DNA. Through evolutionary trial and error, it was found that levers subject to less than one pound of manipulating force preserved the next more detailed level of critical variables better than levers which allowed greater forces to be applied. This chain of reorganizing levels proceeds through an unknown number of steps back to the original reorganizing process that, by active self-induction of mutations, protected against disturbances that could alter the process of replication. That original reorganizing process represented a major step beyond merely resisting such replication-altering disturbances in the primordial soup and evolving through externally-driven mutation and natural selection. Somewhere in there, Genetic Algorithms would fit -- sexual reproduction as another step in developing more and more effective modes of reorganization.

It is the reorganizing system and its ancestors, not the acquired behavioral systems which are rebuilt from scratch in each new lifetime, that constitute the link that extends back through evolutionary history. And throughout this chain of evolving systems, the basic parameter-changing output is RANDOM, so that the environment has no influence on the behavior that controls critical variables. Of course the environment, plus the physical organism as it existed at any time, had a great influence on which internal variables would come to be treated as critical variables. Critical variables are simply those on which, in fact, the continuation of life depends. At no point, however, did the environment have any effect on the most fundamental critical variable of all: accuracy of replication. Controlling that is the point of the whole shebang.

(Penni Sibun (921007.1400), Rick Marken (921007.1230), Bruce Nevin
921008.0646) --

There are some good suggestions for strategy here. I suspect, Rick, that in trying to make the "behavior of perception" make sense to psychologists, you made it make TOO much sense. The shorter the distance between a new way of understanding something and an old one, the easier it is for the listener to conclude that the new one is nothing but the old one. I'm sometimes amazed at the way people can take a perfectly simple and clear explanation, and read all the words so they add up to a totally different explanation. We even have this trouble here, on the net. Every one of us has gone through this sort of misreconstruction of meaning.

I think you may have to go into great detail. You may have to pick out, for each point you make, the nearest conventional idea you can think of and show exactly what the difference is. Did the reviewers get specific enough for you to include their remarks and use them in the paper to spell out the differences? If so, and considering the not-unfriendly reception, this might well be worth doing. You might even consider writing a new paper that deals with just a few of the most important differences. Leading people all the way to what you mean by "behavior IS perception" may involve too big a jump. Maybe you're trying to do in one paper what will really take three or four papers. Pick the size of the wedge to fit the amount that the door is cracked open. This door can be opened only from the inside.

Avery Andrews (several posts) --

I'd like to see that paper on rules and language.

The meeting on simulating societies is a bit premature for us. We really need our Institute for the Study of Living Control Systems, so people can put in the kind of work needed to meet a challenge like this.

Isaac Kurtzer (921007) --

ECS: Elementary Control System. CEV: complex environmental variable.
Both thanks to Martin Taylor.

Bruce Nevin (921007) --

Blindsight seems a little easier to understand in terms of levels of perception. Do things perceived this way always exist as low-level perceptions? That is, would one ever blindly perceive a relationship as opposed to an object?

Best to all, Bill P.

Date: Thu Oct 08, 1992 8:16 am PST
Subject: Living levers, rejection

[From Rick Marken (921008.0830)] Bill Powers (921007.0800) --

The "analogy of the lever" was extraordinary. In three paragraphs or so you were able to make the point I tried to make in my ill- fated "Behavior of perception" paper; and, as usual, you did it far more clearly as well. The best part (for me) -- worthy of considerable reflection -- was:

>From any individual's point of view, ALL interactions with others
>are control of their behavior or influences on their behavior.
>But that is true of all the others too; that's how they see you.
>The people you are trying to influence of control are trying to
>influence or control you, by the very actions and changes of
>properties you see as having been caused by you. To take any one
>side in an interaction is to miss the essence of what is going
>on: interaction.

That one should remain waving, unburned, on every flagpole.

Best regards Rick

Date: Thu Oct 08, 1992 10:32 am PST
Subject: blindsight

[From: Bruce Nevin (Thu 92108 11:48:28)] (Bill Powers (921008)) --

> Bruce Nevin (921007) --
> Blindsight seems a little easier to understand in terms of levels of
> perception. Do things perceived this way always exist as low-level
> perceptions? That is, would one ever blindly perceive a relationship
> as opposed to an object?

Difficult to determine. The article quotes the reports of persons with removed or damaged visual cortices and describes the ability of these people and of similarly lacking monkeys to discriminate visual stimuli that (in the case of humans at least) they cannot consciously see. A flashed light "feels" like a billiard cue advancing or receding. They can discriminate the "texture" perception we have discussed (e.g. gray target vs. B/W grating of equal brightness with lines separated by 3' of arc). Perhaps configurations: for example, one "guessed" at an X vs an O on the basis of "feeling" the "smoothness" of an O or the "jaggedness" of an X; black felt distant, white close; a square was felt as a "sharp movement" or a "corner-shaped wave"; a triangle as "thin waves" and a triangle with rounded sides as "thicker and quicker waves".

A plausible interpretation (which had occurred to me also) is that they are becoming aware of signals that in normal vision are input to configuration detectors, etc., in the cortex, and that normally are ignored, that is, never or rarely come to awareness, because they are after all controlled by those higher-level systems whose outputs in turn are the signals of interest for attention.

The discussion of a sudden regaining vision in adults after some years of training suggests that through reorganization replacements for the missing higher-level systems were developed. An 11-year-old child had virtually normal vision immediately after removal of the occipital lobe (in 1935) because of a prenatal cyst; perhaps the cyst had interfered with vision and he had already been developing parallel systems to get around the problem. (Aside: does it not seem likely that I/O requirements of neighboring higher- and lower-level systems provide a kind of "template" for reorganization to match in creating a new system? Well, I guess obviously: error is not reduced until those neighboring systems are happy. Error at level n+1 and at level n-1 results through reorganization, eventually, in the creation of an ECS at level n. Loss of plasticity due not to running out of neural matter, but rather to having a larger number of possibilities for modification to exhaust on level n+1 and level n-1 before reorg effects reach as far as level n.)

Weiskrantz was recently named a William James Fellow by the APS for his "discovery" of blindsight, inter alia. (I put "discovery" in quotes with thought of his brief review of predecessors who described it but were ignored because it was obviously impossible to see without having a visual cortex. Getting a notion accepted is evidently as prizeworthy as discovery. So persevere, Rick, you might get a Nobel yet!)

Bruce bn@bbn.com

Date: Thu Oct 08, 1992 1:15 pm PST
Subject: Significance & importance

From Greg Williams (921008 - 2) >Bill Powers (921008)

>At no point, however, did the environment have any effect on the most
>fundamental critical variable of all: accuracy of replication.
>Controlling that is the point of the whole shebang.

But you admit that the environment can have great effect on what you consider "surface" phenomena. It appears that you explicitly agree with me that both radical organismism and radical environmentalism are wrong, because the (to you) "significant" variables don't depend on the environment and the (to you) "insignificant" variables do depend on the environment. So, we end up with what I've referred to as organismic/environmental "co-determination" of (some) variables. That's all I've been lobbying for. Of course, you seem to believe that organismism in some sense is vindicated if the (to you) "significant" variables do not depend on the environment. I still think they do, in important ways, but I don't need them to for what I've been calling "co- determination." "Critical variable organismism" -- a limited version of organismism -- is, I think, compatible with the kind of "co-determination" I've been arguing for. Your admitting "co-determination" of (to you) "insignificant" variables is enough for me right now. Even with "critical variable organismism" (or maybe it should be "evolutionism"?), the result is an explicit step away from radical organismism.

This leads to the need to consider why anyone would or would not decide to be interested in having explanations of or, more generally, in understanding and dealing with, how someone's (to Bill) "insignificant" variables can be influenced by that person's environment (both living and non-living parts). Maybe Clark McPhail could help begin to answer this question. Or maybe victims of con artists could help. Or maybe welfare mothers. Or TV watchers. Or drug addicts. Or parents. Or children. Or teachers. Or students. Or counselors. Or prisoners. Or most anyone engaged in and/or studying social interactions.

The resulting data might help to convince some PCTers that the slogan "No one can control you" -- meaning that no one make you want what you don't want -- merits a "Big deal! I can still have plenty of problems -- and plenty of benefits -- due in part to others controlling their OWN perceptions which depend on what I do. And those problems/benefits aren't insignificant TO ME!"

Best,   Greg

Date: Thu Oct 08, 1992 2:33 pm PST
Subject: Re: they already know that

[From Rick Marken (921008.1300)]   penni sibun (921097.1400) --

Ah, it's nice to have you back.

>if you showed how well-known phenomena can be
>analyzed differently and novelly in yr hierarchy--*that* way you can
>make an *argument* for your hierarchy.

Actually I did that, though I didn't do it quantitatively, unfortunately (like by showing a working model) -- mainly because the right data wasn't available and I was not interested in doing the experment at the time -- but I might do it eventually; this was really a theoretical paper, after all. The "well known" phenomenon was really just a finger tapping study done by Rosenbaum (described in JEP:HPP, 1984). He had subjects make sequences of finger taps as rapidly as possible. The result of interest to me was that the speed limit was about 4/sec, the same as the perecptual limit for perceiving auditory and visual sequences. My alternative theory of what was going on (Rosenbaum had an output generation model) was that subjects were controlling a sequences of perceptions of finger tip pressure. I suggested, in the paper, that one piece of evidence to support this would come from a study the subjects' ability to perceive finger pressure sequences that are presented to them (that are not generated by their own taps) -- my guess is the same limit will be found. So the speed of finger tip sequence production is limited by the ability to perceive (and hence control) sequence -- not by limitations in the output process. There are other experiments that I could think of -- the most obvious would be to add disturbances to the finger movements and see that this has little effect on the perceived sequences. But the goal of the paper was to find evidence of perceptual control in the existing literature -- not easy given that so much of the data is fairly useless because it is averaged over people.

>but see, i'm not convinced, from the kinds of things you write about,
>that it's coherent for you to say ``perception is behavior.'' i think
>i can map from your model straightforwardly into one w/ pieces labeled
>``perception'' and ``behavior.'' you haven't convinced me i should
>consider the terms synonymous or what that would mean.

Believe me, I understand your problem. I've been trying to explain and demonstrate this stuff for years and it is not easy to communicate what I consider to be obvious points. I'm sure it's as frustrating for you as for me. I thought I did a pretty good job of explaining in the paper that "behavior" refers to controlled results of action; actions are the means by which behaviors (controlled results) are brought to their intended levels. In a control loop, the result that is untimately controlled is the perceptual input variable. The actions used to produce this results are not "behaviors" from the point of view of the behaving system but they may appear to be "behaviors" from the point of view of the observer of the system, simply because they are produced by the system (the word "behavior" is typically used to refer to any result of an organisms muscle actions). For example, in the "rubber band" demo, the position of the knot is a behavior of the subject; the position of the subject's hand is NOT (though it would be seen as such by any observer of the situation). The position of the knot is controlled -- it is an intentional result of the subject's actions. The position of the hand is not controlled -- it will be moved wherever necessary to keep the knot of target.

I think you raise an interesting point about showing how PCT is a better model of some "well known" phenomenon than some current explanation. One problem we have in PCT is that most of the "well known" phenomena around are, from our point of view, not phenomena at all. We went through this exercise some time ago when we looked, in detail, at a study of voice onset time in phoneme recognition. It turned out that the phenomenon was only there statistically -- at best. Maybe you could pick a "well known" phenomenon, give the reference to the reserach article it is described in (so we can all get a copy of the results) and we can see what PCT has to say about it.

Best regards   Rick

Date: Thu Oct 08, 1992 4:53 pm PST
Subject: Re: Percepts; reorganization; competing behaviors

[Martin Taylor 921008 20:15]   (Bill Powers 921006.1530)

I'm delayed in this response, and I think that I may not be able to respond much over the next 3 or 4 weeks, first because of other pressures and then because I'll be away. I owe Rick a response on statistics based on a private interchange he suggested be continued publicly, and I hope I'll get to that. But here's a little something, anyway.

>RE: competing behaviors
>
>>When one of the mutually competing behaviours acts so as to provide
>>a percept that satisfies the higher-level ECS's reference signal
>>(taking the car gets me closer to my destination), the reference
>>signals for the other would-be competitors are reduced.
>
>What reduces them? I still don't see how you can be riding a bicycle
>and driving a car at the same time.

If the reference for "see myself as riding a bike" is a result of the output from the ECS that is controlling for "see myself arriving at the destination," then nearing the destination will reduce that reference. I suspect, however, that since we are working above the category level, there are threshold effects as well. Be that as it may, when I am at the destination, there is no error in not seeing myself taking the bike. >
>>Then the higher-level reference is not being satisfied, and as you
>>say, the output of the "taking bike" system increases, to the point
>>where it overtakes the "taking car" system and inhibits it.
>
>Where are these outputs acting, outside the system? What kind of
>outputs are they? Come on, Martin, you're waving your arms.

Well, I can't do that and type at the same time, can I?

In the scenario I painted, the output for taking the car was indeed acting. It was acting on the CEV that we call driving the car. But the world made that CEV immovable (the car didn't work, so it, too, was immovable). The output of taking the bike was not capable of causing overt actions in the real world so long as it needed ECSs at lower levels that were acting on behalf of the "taking car" ECS. But when its output became large enough to cause a switch (which I think must occur at the category level and above, there being no intermediate possible reference levels between categories), then those lower ECSs would cease supporting "taking car" and start responding to reference levels that ultimately derived from "taking bike." This doesn't imply internal inhibition. It's just that the relevant CEVs do not admit to simultaneous control in the real world, because in the real world their control requires some lower-level CEVs to be controlled to more than level at the same time, which can't happen.

I don't perceive arm-waving, and I'm certainly not controlling for perceiving it. The linkages and interactions seem quite clear an unobjectionable to me, and I'm not quite sure what it is you are having difficulty with.

I do have to admit that I have not built a simulation.

Martin

Date: Thu Oct 08, 1992 5:13 pm PST
Subject: Re: Why 99%, hierarchical perception

[Martin Taylor 921008 20:30]

In response to Rick Marken's suggestion that we continue a private discussion via CSG-L. In view of his suggestion, I hope he will not object if I quote from that mail. If I'm wrong, I apologise in advance.

I said in a posting just sent that I was going to put this off. But it might be put off for a long time, so rather than do that, I'll try to be brief.

Rick (921005 15:55)

>>Why, for example, am I
>>wrong in suggesting that your data, which can be used (statistically) to
>>decide between x+y and x*y, could be used to discriminate between x*y and
>>log(x)+log(y)?
>
>Why don't we do this on the net. I would have to have a little better
>idea of exactly what you are deciding about (with respect to the model).
>Why, for example, would you expect the variance of the error of prediction
>to suggest anything about nature of the controlled variable? What is
>the statistical analysis that you think is worthwhile for my data?

I'm not sure of the detail of the study (you probably described it somewhere but I have forgotten), but let's assume that subjects were asked to "keep the figure the same size" while you disturbed some aspect of it. You built a simulation model that incorporated appropriate gains and delays, and tried it with a perceptual input function of (x+y) and again with (x*y). You found that the best you could do with (x+y) as an input function left you with 2% unexplained variance, but you could halve that error if you used (x*y). I doesn't matter whether this is exactly what you did, because it can serve to illustrate what I am talking about in any case. What are the follow-on possibilities to this study? What have you found out? I think you might be tempted to believe that you have shown the possibility that perceptual input functions (PIFs) can include multiplication operators, x and y both (presumably) being controllable CEVs in themselves. Now I propose to you that this is false, and I propose as a counter-possibility that addition operators and logarithm operators are possible, and that the "correct" PIF is log(x)+log(y). It is (conceivably) important to distinguish these possibilities, because they may have strong implications for the brain structures involved not only with this percept, but with many other percepts.

What to do? If you have tested only with one "size", you can't tell these possibilities apart very well, since the controlled value of either (x*y) or (log(x)+log(y)) is a constant, and the errors are (by assumption) due to non-infinite gains in some control loops associated with the task, which affect the apparent overall gain of the "size" control loop. So, since the error is a simple difference between reference signal and perceptual signal, all we can do is estimate it and say "that's how good the control is."

But now let's introduce a new reference size, in which x and y are both doubled. So the new level of (x*y) is four times the old reference level. But the gain is, we assume, undisturbed, so we expect to see much the same error in prediction. We have only changed the reference level, an additive variable in the equation. The same is true if the subject is controlling log(x)+log(y), but instead of multiplying the reference signal by 4, we have added 0.6 units to whatever it was beforehand. If we now refer the variance of our estimates to the size set as a reference for the subject, we find that if the logarithmic hypothesis is correct, the residual variance is constant in log size units, whereas if there is a multiplier operator in the PIF, the residual variance is constant in area units.

I think the results would be contaminated by other effects, but the main idea should be valid. It's a very simple kind of case, but it is typical of the sort of thing that psychologists want to know: is shape recognition initiated in the Right Hemisphere and symbolic in the Left, before being transferred across to the other hemisphere for syntactic and pragmatic integration? That's the kind of question of interest for solid practical reasons that apply to cases of stroke. How could you do a non-statistical PCT study to address such a question?

Martin

Date: Thu Oct 08, 1992 7:29 pm PST
Subject: significant (to Bill)

[From Bill Powers (921008.1800)] Greg Williams (921008) --

>But you admit that the environment can have great effect on what
>you consider "surface" phenomena.

Yes. Surface phenomena are those like pushing back against the effect of a disturbance to prevent a controlled variable from changing, or varying an action so as to keep altered properties of the environment from disrupting control. When these things are done in such a way as to keep all an organism's controlled variables still under control, there are no important effects on the organism. All the important effects arise when external influences prevent control from being successful, so controlled variables are no longer under control. Disturbances can have pronounced effects on the way a person acts. But if those actions are successful, and don't prevent other control actions from being successful, such disturbances are insignificant in the life of the organism.

>It appears that you explicitly agree with me that both radical
>organismism and radical environmentalism are wrong, because the (to
>you) "significant" variables don't depend on the environment and >the
(to you) "insignificant" variables do depend on the >environment.

This is a misstatement of my position. My position is that external influences are significant to an organism if they prevent it from controlling its controlled variables. An external influence that does not alter the organism's ability to control any of its controlled variables is insignificant to the organism. It causes no change in organization and it causes no significant error, hierarchical or critical.

>Of course, you seem to believe that organismism in some sense is
>vindicated if the (to you) "significant" variables do not depend on
>the environment.

I say that the organism is controlling successfully if its controlled variables don't depend on external influences. A variable is controlled if it matches its reference level within reasonable tolerances (whether that reference level be changing or constant). What's to "vindicate?" That's just how control systems work, according to PCT. The significant variables (to me) are the controlled variables (significant for one reason) or external variables capable of preventing control from succeeding (significant for quite a different reason).

>I still think they do [depend on the environment], in important ways,
>but I don't need them to for what I've been calling "co-determination."

Are you saying that controlled variables do depend on the environment in important ways? What ways, when control is successful? Note that if a high-level disturbance is countered by a change in a lower-level reference signal, the lower-level perception still tracks the lower- level reference signal (and so remains undisturbed), while the higher- level perception also remains undisturbed (where "undisturbed" means close enough to the reference signal to satisfy all the organisms's purposes and needs). When the hierarchy is operating properly, all perceptions at all levels remain close to their respective reference signals: they remain under control.

>This leads to the need to consider why anyone would or would not
>decide to be interested in having explanations of or, more
>generally, in understanding and dealing with, how someone's (to
>Bill) "insignificant" variables can be influenced by that person's
>environment (both living and non-living parts). Maybe Clark McPhail
>could help begin to answer this question. Or maybe victims of con
>artists could help. Or maybe welfare mothers. Or TV watchers. Or
>drug addicts. Or parents. Or children. Or teachers. Or students. Or
>counselors. Or prisoners. Or most anyone engaged in and/or studying
>social interactions.

Nobody in these categories needs any help in dealing with the world as long as all perceptions are successfully controlled at their reference levels, and all critical variables remain near their reference states. Effects of external events become "significant" only when they frustrate or disrupt or prevent control.

They may also be significant when control has already failed or has never been learned, and external events make control possible again or for the first time. They are never significant when successful control exists.

When you mention these categories, the implication I get is that there is some problem associated with each category. If the welfare mother or the TV watcher is controlling all the variables that the person is organized to control at all the levels that exist in that person, and there is no critical error, the only problem is in the mind of the observer who sees something wrong with what the person is controlling for and would like that person to control for something else. Any attempt to achieve this wish will, of course, result in conflict.

If there is a problem, it is a control problem. There is inner conflict, or interpersonal conflict. Some control system is badly organized and is allowing errors that are large enough to call for reorganization. Some control skill needed to limit critical error has not been learned. The control problem could be in the TV watcher or in the psychologist who thinks that watching TV is bad for the person. Or both. If neither, there is no significant problem.

I hope that this clears up what "significant (to Bill)" means.

>The resulting data might help to convince some PCTers that the
>slogan "No one can control you" -- meaning that no one make you
>want what you don't want -- merits a "Big deal! I can still have
>plenty of problems -- and plenty of benefits -- due in part to
>others controlling their OWN perceptions which depend on what I do.
>And those problems/benefits aren't insignificant TO ME!"

Problems are problems only if you want to fix them and can't. Benefits are problems only if you can't get them when you want them. If you have plenty of problems, they are control problems. They will be solved by solving the control problems. The mechanisms for arriving at such solutions are inside each organism, not in the environment between them. Even the mechanism for understanding, accepting, and putting into practice a solution communicated by someone else lies inside the recipient. This is the basis on which I would use control theory to approach a person with such problems -- not by looking for ways to change the person's environment (although that is not ruled out) but by helping to eliminate conflict and enhance the person's ability to regain control. This would apply to everyone involved in the problem.

I would NOT explain to the person that the problems are co-determined by the person and by the environment, even though that does pretty much cover the possibilities. Even if the problem is partly determined by the environment, the person is going to have to get that aspect of the environment under control, as perceived, of course, in order to do anything about it.

Best,   Bill P.

Date: Fri Oct 09, 1992 3:13 am PST
Subject: An ounce of prevention...

From Greg Williams (921009)   >Bill Powers (921008.1800)

>Disturbances can have pronounced effects on the way a person acts. But
>if those actions are successful, and don't prevent other control
>actions from being successful, such disturbances are insignificant in
>the life of the organism.

Do you disagree that the actions taken by a person at time t1 which are necessary for successful control by that person at time t1 could result in unsuccessful control (and hence Bill-significance) at time t2? An example is the successfully controlling person paying for desired expensive shoes (t1) with his/her credit card and later (t2) having to deal with not getting to go on a desired trip to the Riviera because of not having enough money.

Do you disagree that the actions taken by a person at time t1 which are necessary for successful control by that person at time t1 could result in improved control (and hence, let us say, Greg-significance) of some percepts (possibly distantly related to the percepts being successfully controlled at time t1)? An example is putting away paycheck savings (t1) and after retirement (t2) using them to pay for a trip to the Riviera.

>My position is that external influences are significant to an organism if
>they prevent it from controlling its controlled variables. An external

>influence that does not alter the organism's ability to control any of its
>controlled variables is insignificant to the organism. It causes no change in
organization and it causes no significant error, hierarchical or critical.

In the first example above, if the ("kindly") storekeeper refused to take the person's credit card (t1), judging (quite patronizingly) that this particular assemblyline worker really shouldn't buy a $400 pair of shoes, then that "external influence" would be Bill-significant at t1 and Greg-significant at t2, wouldn't it? In the second example above, if the savings bank folded (no FDIC!) before t2, then that "external influence" would be Bill-significant at t2, wouldn't it?

>The significant variables (to me) are the controlled variables
>(significant for one reason) or external variables capable of
>preventing control from succeeding (significant for quite a different
>reason).

To which I add: external variables capable of facilitating successful control (allowing it where otherwise it would be impossible) -- significant for a third quite different reason.

>>I still think they do [depend on the environment], in important
>>ways, but I don't need them to for what I've been calling "co-
>>determination."

>Are you saying that controlled variables do depend on the environment
>in important ways? What ways, when control is successful?

Historically. Environmental influences prior to time t1 can affect the trajectory of controlled variables after time t1. This is the learning/ reorganization (my "long-term influence") disagreement we've been having.

>When the hierarchy is operating properly, all
>perceptions at all levels remain close to their respective reference
>signals: they remain under control.

But the hierarchy (or, more generally, control organization) CHANGES over time. You think those changes aren't due in any way to the environment; I think the environment plays an important role.

>Nobody in these categories needs any help in dealing with the world as
>long as all perceptions are successfully controlled at their reference
>levels, and all critical variables remain near their reference states.
>Effects of external events become "significant" only when they
>frustrate or disrupt or prevent control.

I wonder if you practice automotive "preventive maintenance"? Suffice it to say that there are many people who (patronizingly or not) believe in keeping drunks from behind wheels, and so forth, to attempt to PREVENT the disruption of control by people. Are you REALLY pleading for a completely "hands-off" approach to regulating social interactions? Should the police wait until AFTER the wreck to arrest the drunk? When the victim is dead, it is difficult to give him/her any help with controlling!

>When you mention these categories, the implication I get is that there
>is some problem associated with each category. If the welfare mother
>or the TV watcher is controlling all the variables that the person is
>organized to control at all the levels that exist in that person, and
>there is no critical error, the only problem is in the mind of the
>observer who sees something wrong with what the person is controlling
>for and would like that person to control for something else. Any
>attempt to achieve this wish will, of course, result in conflict.

Is there a problem in the mind of an observer who predicts that the presently successfully controlling TV watcher is going to have problems controlling in the future which are due in part to his/her current ACTIONS and attempts to alter how the person ACTS, rather than what the person is controlling for?

>>The resulting data might help to convince some PCTers that the
>>slogan "No one can control you" -- meaning that no one make you

>>want what you don't want -- merits a "Big deal! I can still have
>>plenty of problems -- and plenty of benefits -- due in part to
>>others controlling their OWN perceptions which depend on what I do.
>>And those problems/benefits aren't insignificant TO ME!"

>Problems are problems only if you want to fix them and can't.

For sure. And you certainly can't fix getting killed by a drunk driver after you're already dead. So you attempt to alter the drunk's ACTIONS beforehand -- you don't try to alter his/her belief than drunk-driving is a God-given right! -- and deliberately conflict his/her control system by throwing him/her in jail. Then you give him/her the choice of alcoholism treatment or no driver's license.

>Benefits are problems only if you can't get them when you want them.

For sure. And if you come to want them when it's too late to get them, perhaps because a con artist has successfully controlled for certain of his/her perceptions dependent on certain of your actions, or perhaps because your teacher was a lousy communicator, or perhaps because your parents believed that you should do whatever you want to do and not what they think is best, too bad!

>I would NOT explain to the person that the problems are co-determined
>by the person and by the environment, even though that does pretty
>much cover the possibilities. Even if the problem is partly determined
>by the environment, the person is going to have to get that aspect of
>the environment under control, as perceived, of course, in order to do
>anything about it.

Do you prefer telling the widow of the drunk driver's victim that she's just going to have to regain control of her perceptions now that he husband is gone to telling the drunk to get into the police car OR ELSE?

Greg

Date: Fri Oct 09, 1992 6:18 am PST
Subject: significance of the insignificant

[From: Bruce Nevin (Fri 92109 09:30:09)]   (Greg Williams (921008 - 2)) --

>   . . . how someone's (to Bill) "insignificant" variables can be
>influenced by that person's environment (both living and non-living parts).
>Maybe Clark McPhail could help begin to answer this question. Or maybe victims
>of con artists could help. Or maybe welfare mothers. Or TV watchers. Or drug
>addicts. Or parents. Or children. Or teachers. Or students. Or counselors. Or prisoners.
Or most anyone engaged in and/or studying social interactions.

I have only delved a bit into the Collected Writings of Milton Erikson, but enough to suggest to me that you will find a lot of grist for this mill there.

One anecdote about Erikson: He's giving a lecture. At a certain point two people, from opposite sides of the auditorium rise, walk up on the stage, and sit down in chairs near the podium. When asked why they did so, they are somewhat puzzled, and say it just seemed like the right thing to do. He then explains that he communicated with them individually during the course of the lecture delivered to the whole audience. Catching the eye of this one while saying a word or phrase that had a general meaning for the audience, but a specific one for the person. Gesturing in a way that had general meaning for the audience but that indicated what was desired to the individual on that side. Embedding suggestions covertly within the flow of communications that had different overt purposes. In this he was making use of what he called "everyday trance," a state that is not limited to people listening to lectures. This demonstration then having its purpose to exemplify points made in the lecture, he thanked his still somewhat bemused assistants for their help, in a way that they felt pleased and not exploited, and concluded the lecture.

The collected papers are in four volumes, edited by a student of his named Rossi.

I will not defend Erikson or his work. He speaks quite well for himself. But you do have to read the material. Many of his approaches seem to me to rest on creative use of ambiguity, rather than on drawing attention to lower levels of perception (e.g. the stereotype Bill mentioned of the voice droning on about heavy eyelids), but that may have the same effect (e.g. his handshake technique).

Bruce bn@bbn.com

Date: Fri Oct 09, 1992 11:28 am PST
Subject: skinner person

possibly the person i have tried to reach - the one who posted about an upcoming article clarifying skinnerian behaviorism - has not been on the net in the past few days, so now ii will repeat the mesaage :

1) i AM interested in a copy of the paper (a pre-print draft is fine)
2) please send to isaac kurtzer p.o. box 4509 SFASU Nacogdoches,TX 75962
3) please give me an indication either way thankyou

isaac n. kurtzer

p.s. to brother bourbon - please give me a short list on "the greatest hits" of skinner et. al.

p.p.s. to rick "the player" marken - personally i feel you shouldn't try to take your paper through the back door; there is simply no way to reconcile PCT with contemporary psychology since a it hits at the assumptions of psychology which are inarguable and other reasons basically Kuhnian in nature- i'm sure you get the drift. i say wait till their nil-point palace crumbles!!!! i would like a copy of your paper,please.

Date: Fri Oct 09, 1992 12:29 pm PST
Subject: Ounce of prevention ...

[From Bill Powers (921009.0800)] Greg Williams (921009) --

>Do you disagree that the actions taken by a person at time t1 which
>are necessary for successful control by that person at time t1
>could result in unsuccessful control (and hence Bill-significance)
>at time t2?

I don't disagree. In that case, the action DOES affect something else that the person is controlling for, and creates loss of control. After a person has done this for a few years, presumably some reorganization will take place and the person will learn to control for predicted perceptions. Of course if you prevent the person from doing the improvident act, no learning will take place.

>Do you disagree that the actions taken by a person at time t1 which
>are necessary for successful control by that person at time t1
>could result in improved control [at time t2]...?

Nope, don't disagree with that, either. This kind of control would belong at about the program level, wouldn't it? Or perhaps principles (saving for the future). If the improvement at time t2 were not intended, however, it would just be a lucky accident, and could have turned out oppositely.

>In the first example above, if the ("kindly") storekeeper refused to take
>the person's credit card (t1), judging (quite patronizingly) that this
>particular assemblyline worker really shouldn't buy a $400 pair of shoes,
>then that "external influence" would be Bill- significant at t1 and
>Greg-significant at t2, wouldn't it?

If the customer doesn't reorganize, he/she will either protest loudly and violently at the discrimination or find another store that will sell the shoes. If that doesn't work, the customer still won't control for saving for the future; evidently there's no reference level for doing that, and a storekeeper or storekeeper conspiracy can't install one. So the customer will blow the money on something else. It's very unlikely that by the time t2 arrives, there will be any money left for the trip, no matter what the storekeeper does. That's not what the customer is controlling for.

>To which I add: external variables capable of facilitating
>successful control (allowing it where otherwise it would be
>impossible) -- significant for a third quite different reason.

You can't facilitate a control system that doesn't exist. If the control system does exist and isn't succeeding, the person is probably reorganizing and will eventually succeed or change the goal. If the control system is succeeding, then any direct help will be resisted. If you facilitate control by making the environment more amenable to control, then reorganization will stop and the struggling control system will succeed. The next time the situation comes up, you will probably get a phone call asking if you would kindly set up the environment that way again, please. The third time, the phone message will be "Ready whenever you are." The fourth time it will be "I'll be needing you at 10:00, and don't be late." All you've done is to make yourself part of someone else's control loop.

>>Are you saying that controlled variables do depend on the
>>environment in important ways? What ways, when control is
>>successful?

>Historically. Environmental influences prior to time t1 can affect
>the trajectory of controlled variables after time t1.

I disagree totally. The trajectories of controlled variables are the trajectories of their reference signals, not of environmental disturbances of those variables. The whole point of controlling a variable is to render it independent of external influences and make it depend only on the internal desire that it be in a particular state. If you want to wash a cup, its previous history makes no difference at all. It gets washed, even if it's clean, and whether you pick it up or someone tosses it to you.

>>When the hierarchy is operating properly, all perceptions at all
>>levels remain close to their respective reference signals: they
>>remain under control.

>But the hierarchy (or, more generally, control organization)
>CHANGES over time. You think those changes aren't due in any way to
>the environment; I think the environment plays an important role.

"Due in any way to the environment" is backing down quite a lot, isn't it? Obviously the changes would be drastically affected by the environment if its mean temperature dropped to -100 c. If that's the kind of thing you want me to "admit", then OK, no problem. The organism has to learn to control its perceptions in the world that exists. It has to find one of the actions that will have the required effects on a perception in order to control it. It even has to find perceptions that are (a) controllable, and (b) relevant to maintaining critical variables at their reference levels. All this requires that the properties of and events in the environment be taken into account. If the only eating implements available are a knife and a fork, and the person thinks it impolite to eat with the fingers, the soup will be eaten rather slowly or the person will say to hell with it and slurp out of the soup-bowl. That's because of the environment. But the environment didn't tell the person to eat the soup.

This is nothing that I haven't said dozens of times, but you seem to pay no attention and you put words in my mouth to the effect that the environment doesn't play any role at all. It does play a role: it's there and it behaves according to its properties, according to what you do to it. It does things all by itself, which is how disturbances arise. If the environment changes in some basic way, the organism has to reorganize so as to maintain control of it, and the new organization will one of those that compensates for the change.

My whole point is that none of the changes in the organism would occur if the organism itself didn't actively make those changes, for its own purposes.

>I wonder if you practice automotive "preventive maintenance"?

Yes, some of my perceptions actually encompass an imagined future, which I control. Of course after I've faithfully changed the oil 20 times, the water pump fails. The imagined future is only imagined. Do you do preventive maintenance on your water pump?

>Suffice it to say that there are many people who (patronizingly or
>not) believe in keeping drunks from behind wheels, and so forth, to
>attempt to PREVENT the disruption of control by people.

Fine, they're controlling their own perceptions, aren't they? You may have quite a struggle in getting the car keys away from the drunk, but being sober and I hope bigger, you will win the conflict and have it your way. The next time, the drunk will have a spare set of keys in his sock.

Now you're getting into what is OBJECTIVELY good for people. This is how most major conflicts between people are created.

>Are you REALLY pleading for a completely "hands-off" approach to
>regulating social interactions? Should the police wait until AFTER
>the wreck to arrest the drunk?

I seem to recall saying that if my kid runs into traffic, I grab the kid and get it out of the street, and to hell with avoiding conflict. You do what you think is right. We all do, including Saddam Hussein. If you strongly want people to behave in a certain way, you make laws and establish a police FORCE. You do whatever it takes to get what you want, as long as all your goals are satisfied, including your goals for how people should relate to each other and what kind of society you want to live in. I don't know where you got the idea that I'm pleading for a completely hands-off approach to social interactions.

I'm only pointing out who is responsible for what. If you decide to drag the drunk out of the car, that's your decision carried out to fit your own goals. The environment didn't make you do that; the drunk didn't make you do that. You did it because of what you believe, what you predict, and what you want to happen or not happen; that's your responsibility. You did it as part of making the world be the way you want to perceive it. You can, of course, argue that the drunk and his or her family is better off for what you did; that argument, too, is your responsibility and reflects your convictions and goals. You can't disguise your responsibility by saying "Yes, but he REALLY would have killed himself and orphaned his children!" That's how you understand it. It's your goal that he not be killed and that his children not be orphaned and his wife widowed. Your goal, inside you. You can use logic, statistics, persuasion, ridicule, hypnotism, violence, appeals to right and wrong, quantum mechanics, or whatever tools you command to achieve your goal and to get others to seek the same goal, and it will still be your own goal and yours alone. You can't say "the environment made me do it." The environment is just what it is.

>Is there a problem in the mind of an observer who predicts that the
>presently successfully controlling TV watcher is going to have
>problems controlling in the future which are due in part to his/her
>current ACTIONS and attempts to alter how the person ACTS, rather
>than what the person is controlling for?

Not if this control process succeeds. But it's still the controller arranging things to suit himself or herself on the basis of private beliefs and goals. Unfortunately, it will be difficult to change how the TV watcher ACTS without disrupting the perception under control, namely, seeing the program on the TV. The result will probably be conflict, which may create problems. The conflict is that the watcher wants to be watching TV, and the controller wants the watcher NOT to be watching TV. If conflict is to be avoided, the watcher has to change the goal, not the ACTION of sitting down and turning on the TV and directing the bloodshot eyes to it. The controller can't change the action without disturbing the controlled variable. At best, the controller can try to create a conflict and hope it will be resolved as the controller wants: "OK, I'm driving into town to see the circus, anybody want to come along?" Most likely, the dedicated controller will walk up to the TV set, turn it off, and announce "That's all for today." Implying, "and there's nothing you can do about it, kid."

Except go to a friend's house where the TV is on 24 hours a day.

>And you certainly can't fix getting killed by a drunk
>driver after you're already dead. So you attempt to alter the
>drunk's ACTIONS beforehand -- you don't try to alter his/her belief
>than drunk-driving is a God-given right! -- and deliberately
>conflict his/her control system by throwing him/her in jail. Then
>you give him/her the choice of alcoholism treatment or no driver's

>license.

Right. Overwhelming physical force always works as a way of making other people satisfy your goals, if you have the necessary resources and are willing to watch your back from then on.

>...if you ... want [benefits] when it's too late to get
>them, perhaps because a con artist has successfully controlled for
>certain of his/her perceptions dependent on certain of your
>actions, or perhaps because your teacher was a lousy communicator,
>or perhaps because your parents believed that you should do
>whatever you want to do and not what they think is best, too bad!

Yep, too bad. You can't control what you're not controlling for. That's how it works. If you don't learn from experience you're going to get conned over and over. If you had a bad education it's going to stay bad until you decide to continue it. If your parents let you do anything you wanted, you're going to want a lot a weird things unless you learn from experience that this doesn't get you what you want. If your parents make you do what they think is best, and it's not best, you're still screwed. That's life.

>Do you prefer telling the widow of the drunk driver's victim that
>she's just going to have to regain control of her perceptions now
>that he husband is gone to telling the drunk to get into the police
>car OR ELSE?

Of course not. I'd try to keep the drunk out of the car and see that he got safely home. I might try to steer him to AA on the way. There are lots of things I would try to do because of my own goals for the kind of world I want and the kind of person I want to be in it. If the drunk did kill himself anyway, which is likely, I'd tell the widow that he is really dead and she must adjust her perceptions to fit that fact, and make sure her kids know how he killed himself.

Control theory fits the world as it is, and people the way they are. But an understanding of control theory leads to assigning causality differently, and as a result will change the way people do things and the reasons for which they do them. A resort to arbitrary control of others by means that rest ultimately on violence will be seen as a defeat, not an accomplishment. A viable steady-state society clearly can't be organized on those principles, even though they work in the short term and are all we have available now. Even if you use them and think they're working just fine. When everyone realizes this, a different organization of society will result. If it doesn't, I'll kill everybody, and THAT will teach them.

Best,   Bill P.

Date: Fri Oct 09, 1992 2:06 pm PST
Subject: Re: Why 99%, hierarchical perception

[From Rick Marken (921009.1230)]

Martin Taylor (921008 20:30)

>In response to Rick Marken's suggestion that we
>continue a private discussion via CSG-L. In view of his suggestion, I hope
>he will not object if I quote from that mail. If I'm wrong, I apologise in
>advance.

I don't object, I rejoice (well, anyway, yes, that's what I wanted).

>I think you might be tempted to believe that you have shown the possibility
>that perceptual input functions (PIFs) can include multiplication operators,
>x and y both (presumably) being controllable CEVs in themselves. Now I
>propose to you that this is false, and I propose as a counter-possibility
>that addition operators and logarithm operators are possible, and that
>the "correct" PIF is log(x)+log(y).

No problemo. I was just trying to show that the variable being controlled was more like x*y than x+y; I didn't care how the neural signal was derived from the sensory inputs, though I think it might be interesting to try to figure it out -- especially if you are interested in how neural processing works.

>It is (conceivably) important to distinguish these
>possibilities, because they may have strong implications for the brain
>structures involved not only with this percept, but with many other percepts.

OK.

>What to do?

Probabaly try to find the neural signal that corresponds to the controlled variable and start dropping some single cell recorders in there.

>But now let's introduce a new reference size, in which x and y are both
>doubled. So the new level of (x*y) is four times the old reference level.
>But the gain is, we assume, undisturbed, so we expect to see much the same
>error in prediction. We have only changed the reference level, an additive
>variable in the equation. The same is true if the subject is controlling
>log(x)+log(y), but instead of multiplying the reference signal by 4, we have
>added 0.6 units to whatever it was beforehand. If we now refer the variance
>of our estimates to the size set as a reference for the subject, we find that
>if the logarithmic hypothesis is correct, the residual variance is constant
>in log size units, whereas if there is a multiplier operator in the PIF, the
>residual variance is constant in area units.

I'm not sure I buy this derivation. What is the "variance of our estimates" by the way? What is the "residual variance" here? the variance of the the size of the quadrilateral relative to the reference size? In what way is the residual variance "constant" in log size or area units. I don't understand the analysis, I suppose. If you explain it to me I could do it in a second since it's easy as pie to do the experiment -- I already have the area control program ready to go and it's pretty easy to change the size of the reference square. Maybe it's my bias but why not have the reference square vary continuously in size? Then you could look at the size of the error as a function of the size of the reference area. Maybe that's what you want anyway -- maybe you are saying that the size of the error will be constant with respect to log size rather than size?? Is that it? I still don't get WHY this would be true if their is a log transform applied to the inputs to the perceptual function. Maybe you could give a simple mathematcal reason for this prediciton -- but keep it simple; I'm certainly no math whiz, as, I imagine, you can tell.

>I think the results would be contaminated by other effects, but the main idea
>should be valid. It's a very simple kind of case, but it is typical of the
>sort of thing that psychologists want to know.

Yes, and they miss the whole point. The fact that the person is controlling a variable and that it took some doing to figure out what it is just goes right by. Early astronomers wanted to know how the patterns they saw in the night sky were related to one's personality. I know it's rude but, frankly, I could care less what conventional psychologists want to know. What they want to know is based on their preconceptions about how people work; so they think it is important to know how reinforcement affects behavior or which behavioral/cognitive capabilities are in the right brain and which are in the left, etc etc. Since their preconceptions about how people work are wrong (because they don't understand the consequences of the fact that organisms are locked in a negative feedback situation with respect to their environment -- cf. the blind men paper) it is only by chance that anything psychologists might want to know about is anything more than an illusion from the point of view of PCT.

Have a nice weekend all   Rick

Date: Fri Oct 09, 1992 2:53 pm PST
Subject: Inquiry: Skinner Paper

Subject: RE: skinner person   [From: Dennis Delprato]

(isaac kurtzer)

What you requested is in U.S. mail. And I indicated "OK" re. responding to questions or some such thing.

Dept. of Psychology
Eastern Mich. Univ.
Ypsilanti, MI 48197
Psy_Delprato@emunix.emich.edu

Date: Fri Oct 09, 1992 3:39 pm PST
Subject: Re: Ounce of prevention (hidden bombs)

[Martin Taylor 921009 17:15] (Bill Powers 921009 0800)

>Yep, too bad. You can't control what you're not controlling for.
>That's how it works. If you don't learn from experience you're going
>to get conned over and over. If you had a bad education it's going to
>stay bad until you decide to continue it. If your parents let you do
>anything you wanted, you're going to want a lot a weird things unless
>you learn from experience that this doesn't get you what you want.

More technically, but I think saying the same thing, you are talking about the bomb in the hierarchy that I tried to introduce a few weeks ago. In other words the "school of hard knocks" is not a bad school if you expect to live in an environment of hard knocks. If you can guarantee living where slaves wait on your every desire, you may be able to get the weird things you want whenever you want them. But most of us can't do that. If we don't learn what gets us what we want in a variable world, we won't get it when we are out in that world, because we haven't reorganized in such a way that our control loops have high negative gain through many parallel optional feedback paths (micro-loops) through the world.

The feedback loop from the output of a high-level ECS back to its perceptual signal goes through the real world (or through imagination) by way of many parallel paths I call micro-loops. Even when the overall feedback gain is nicely negative, it is by no means guaranteed that each of the micro-loops also would provide negative feedback on its own. When the real world blocks the error-correcting action of micro-loops that are normally useful or are anticipated to be useful, then it can happen that the overall feedback gain is reduced and can even go positive. The error "blows up" because the person does something that not only doesn't get what he or she wants, but actually gets something distinctly unwanted. Such bombs can lie hidden for a long time in a seldom disturbed hierarchy. When I introduced the bomb concept, I speculated that the ability of bombs to stay hidden in "coddled" nets might account for why so many mass murderers are characterized as Quiet, Nice Kids Who Never Were Any Trouble To Their Parents. Losing one's temper and doing something maladaptive as a consequence of frustration reveals the bomb. Wanting "a lot of weird things ... [that don't] ... get you what you want" is the content of the bomb.

Martin

Date: Fri Oct 09, 1992 3:40 pm PST
Subject: Re: MIT talk

[Martin Taylor 921009 17:45] (Bruce Nevin 921002 13:14:03)

>The theory under discussion holds that the object selected by the vision
>system will be the least complex of the available alternatives.
>Experimental data supporting the theory will be reported.
>
>This work is based on the pioneering ideas of Solomonoff and Kolmogorov,
>and on the more recent ``minimum description length'' concept of Rissanen.

Interesting. In the mid-70's I was very concerned with exactly such concepts, and it is a good part of the reason why I harp so much on statistics as an essential part of perception. I also used the concept of "minimum description length" as an argument for why Occam's razor was not only a handy guide, but a nearly infallible rule for favouring one theory over another, as well as for why we perceive differently using active exploration as opposed to being presented passively with "sensory stuff." But I didn't know about Kolmogorov at that time, nor that it was a novel idea. It just seemed obvious, not really publishable. It still does, and I'm rather surprised that people can make seminars out of it.

There's also an article on this in the (most?) recent Science, which I don't have with me so I can't check whether it is the same authorship (Thomas Marill).

Martin

Date: Fri Oct 09, 1992 3:48 pm PST
Subject: Re: Why 99%?

[Martin Taylor 921009 17:50] (trying a quick backlog clearing operation)
(Tom Bourbon 921002 -- 10:50)

>>Martin Taylor 920929 16:00
>
>>My presumption is that you get the 99% prediction because the subsystems
>>(perhaps ECSs) that are involved in the task are those that support
>>very many different kinds of behavior, and so are not readily disturbed
>>by contextual differences.
>
>What do you mean by "contextual differences?" I believe your presumption
>is wrong, but I am not certain what you are saying. By "contextual
>differences" do you mean that target positions follow different random
>paths on every trial and the cursor is disturbed by a different random
>function on every trial?

No, not at all. I mean that a pianist uses the same muscle tension control systems to create a legato run as to pick up a piece of bread or to wave to a friend. But the higher-level system that generates (eventually) the reference level changes that direct the muscel tension control systems can be affected by the pianist's hearing a flute played in the next apartment. That could disrupt the legato run. The lower-level systems, as they support many widely different kinds of perceptual control at higher levels, are not very susceptible to contextual influences, whereas the higher level ones, being more specific, can be more disturbed.

When you model the pianist, you will probably get very good prediction for the low-level controls, but when the flute disrupts the legato run, your simulation will not model it. You could build in that source of disturbance, but the odds are that you wouldn't.

Martin

Date: Fri Oct 09, 1992 3:49 pm PST
Subject: Re: reorganization

[Martin Taylor 921009 17:00] (Bill Powers 921006.1530)

> Just by trial and error, I found that the most
>reliable method (so far) entails (roughly) making the rate of
>reorganization depend on the rate of change of absolute critical error
>TIMES the absolute critical error, so that the changes get smaller as
>the remaining error diminishes. I realized only afterward that this
>amounts to taking the derivative of the SQUARE of the error, which
>neatly takes care of assuring positive values and that the result will
>give the best least-squares fit to the optimum solution. You might try
>this in addition to your (e*e + k*e*de)/G.
>In other words, try d(e^2)/G.

Did you try varying the relative contributions of the value of e**2 and of d(e**2)/dt? In other words, varying k in my expression between zero and infinity? My formula is yours if k is large, but if k is moderate, then reorganization could occur with a sustained error that does not change value, which seems as if it would be a good thing to do.

>Another tip. Rather than just make random corrections based on

>critical error, I have found that it's best to randomly change a
>parameter that determines the direction and speed of changes in
>parameters. Between reorganizations, the changes then continue on each
>iteration. This is strictly analogous to the E. coli method of
>locomotion, in which, between tumbles, the organism continues to swim
>in a straight line, continuously altering its relationship to a radial
>gradient.

This sounds as though you are doing Hebbian learning (on the output link weights?) rather than what I characterize as "reorganization" (the changing of link sign or of the link targets). I can't see what this paragraph would mean when applied to output links that had weights of only +-1 or zero. I thought that the simple control nets you were working with all had this restriction on the output weights, and that you were waiting to be forced into allowing real-valued output weights.

My intuition has been that you would be forced into using real-valued output weights when you got into massively parallel control nets in which no single ECS could be responsible for the whole of any behaviour and there was permanent tension in the net. But so far as I am aware, you haven't studied that kind of net yet. Am I wrong? Or did you get into real-valued weights when you found that sign-flipping had nasty effects? Or are you talking about perceptual input functions here?

>>In the Little Baby project, the second (probably) experiment will >be
>>to start with a random set of output connections, and use Hebbian
>>learning on the perceptual input functions to see whether the baby
>>can learn to perceive the world in a way compatible with its (fixed
>>randomly assigned) outputs.
>
>I can tell you already that this will work, using my method of
>reorganizing outlined above, with up to 20 independent control systems
>controlling through 20 shared environmental variables. The critical
>variable is total squared error across all systems. Reorganization is
>applied globally, not based on each system's error.

Again, I'm not 100% sure what you were varying. I had assumed that your reorganization experiments were restricted to changing the output system rather than the perceptual input system. If you have studied error-induced modifications of the perceptual input functions, we would love to know more about your results. We don't mind re-inventing wheels, but it is nice to know that our wheels work like yours if we do. (By "we" I mean Chris and me, though I expect the CSG-L readership would be interested, too).

Martin

Date: Fri Oct 09, 1992 3:51 pm PST
Subject: Re: Misc subjects

(Martin Taylor 921009 1800] (Bill Powers 921002.0600)

On The Test

>RE: applying disturbances (stuck phonograph record division):
....

>The disturber should NOT control for a visible change in the variable
>being disturbed. That simply creates conflict. What the disturber
>should do (where possible) is alter some OTHER variable that is
>loosely coupled to the putative controlled variable. This will elicit
>an opposing change in the controller's actions even if the controlled
>variable doesn't visibly change. In low-gain situations (like the coin
>game) this doesn't matter so much. But when the control system
>involved is a very good one, insistence on seeing the controlled
>variable actually change will result in applying very large forces to
>the controlled variable, with a probable change in what variables the
>controller is controlling. The Test is most accurate when the
>controlled variable doesn't change at all (that you can see).

I wouldn't have thought it made any difference whether the observer affected the supposed controlled CEV directly or through a "loosely coupled variable." And I wouldn't have thought that actually seeing a visible effect of the disturbance on the test CEV was important. If you do, the subject probably isn't controlling it with any substantial gain. What the observer should observe is a failure of the test CEV to change as much as it should have been expected to do, given the magnitude of the applied disturbance. If the observer wants to discover the limits of the subject's control, then what you talk about will presumably come to pass--large forces and a change in the subject's behaviour (if not organization).

Martin

Date: Sat Oct 10, 1992 4:29 am PST
Subject: A POUND of prevention...

From Greg Williams (921010) >Bill Powers (921009.0800)

GW>>To which I add: external variables capable of facilitating
GW>>successful control (allowing it where otherwise it would be
GW>>impossible) -- significant for a third quite different reason.

>You can't facilitate a control system that doesn't exist. If the
>control system does exist and isn't succeeding, the person is probably
>reorganizing and will eventually succeed or change the goal. If the
>control system is succeeding, then any direct help will be resisted.

I was thinking of the case where the person didn't even CONSIDER controlling some percepts until AFTER the facilitation. Like taking the con artist to court AFTER the policeman informs you that you've been conned (you didn't know that before) and that the con artist is currently in jail on another charge; if you hadn't been told that you were conned, you wouldn't want to control for the perception of taking the con artist to court, and such control would be impossible if the con artist hadn't been detained by the police. I don't see reorganization going on here, just the same hierarchy operating successfully with new information.

BP>>>Are you saying that controlled variables do depend on the
BP>>>environment in important ways? What ways, when control is successful?

GW>>Historically. Environmental influences prior to time t1 can affect
GW>>the trajectory of controlled variables after time t1.

>I disagree totally. The trajectories of controlled variables are the
>trajectories of their reference signals, not of environmental
>disturbances of those variables. The whole point of controlling a
>variable is to render it independent of external influences and make
>it depend only on the internal desire that it be in a particular
>state. If you want to wash a cup, its previous history makes no
>difference at all. It gets washed, even if it's clean, and whether you
>pick it up or someone tosses it to you.

I'm not talking about any current environmental dependence of controlled variables in the currently existing control structure. I'm saying that the control structure itself at time t1 depends in part on the environment prior to t1. (1) WHETHER learning/reorganization takes place depends in part on the environment (I think we agree on this). (2) WHERE learning/reorganization stops (the resulting new control structure) depends in part on the environment (you don't agree with this). (3) Attempts to alter another's current reference signals "directly" (not involving learning/reorganization) is likely to produce conflict in the current control structure (I think we agree on this).

>My whole point is that none of the changes in the organism would occur if
>the organism itself didn't actively make those changes, for its own purposes.

I accept your point. But it is only part of the story, since the changes would be different if the environment were different. There is a joint contribution to the trajectory of control from both organism and environment. That the contribution from the organism is purposive and the contribution from the environment is not purposive is true only when the environment does not contain other controllers.

>But an understanding of control theory leads to assigning causality
>differently, and as a result will change the way people do things and
>the reasons for which they do them.

I don't think a difference in assigned causality will change much.

>A resort to arbitrary control of others by means that rest ultimately on
>violence will be seen as a defeat, not an accomplishment.

I do think that a PCT-understanding of the ways in which control of perceptions depending on others' actions must NOT be arbitrary if conflict and probable violence are to be avoided -- that is, an explanation of the CONSTRAINTS on non-conflicting control of others' actions -- might result in what my ideology would count as a better world. And that sort of understanding/explanation is what I'm interested in coming up with. People "the way they are" are controllers. This fact isn't going to go away. Control of others' actions isn't going to go away. But SOME conflicting/violent control might, given sufficient understanding that there ARE ways to see others act the way you want them to act WITHOUT using force or the threat of force. Understanding that it is ultimately problematic for YOU to try to alter (in the short-term) another's wants is only the first step; the next step, given that you MUST control some perceptions which depend on others' actions, is understanding how non-conflicting control of others' actions is possible, and understanding how it must be done. The only alternative is to try to convince you to STOP trying to control ANY of your perceptions which depend on others' actions. Do you think that alternative can work?

Best,   Greg

Date: Sat Oct 10, 1992 6:58 am PST
Subject: No reorganization

[From Bill Powers (921010.0900)]   Greg Williams (921010) --

>I was thinking of the case where the person didn't even CONSIDER controlling
>some percepts until AFTER the facilitation. Like taking the con artist to
>court AFTER the policeman informs you that you've been conned (you didn't
>know that before) and that the con artist is currently in jail on another
>charge; if you hadn't been told that you were conned, you wouldn't want to
>control for the perception of taking the con artist to court, and such
>control would be impossible if the con artist hadn't been detained by the
>police. I don't see reorganization going on here, just the same hierarchy
>operating successfully with new information.

Dear Mr. Williams,

The Lexington District Attorney's office wishes to inform you that William T. Powers, also known as L. Ron Hubbard III., is presently being held by Lexington County police on a charges of extortion, theft of services, and fraud in connection with his promotions of "dianetics," Rosicrucian therapy, and stock manipulation. He has also been promoting a new scheme called "control theory," telling old people that using this theory they can extend their lives and increase their income, and accepting large "donations" from them. In support of this con game, he and a woman posing as his wife have persuaded several small publishers to provide books at very low cost, which they give to these old people as proof of legitimacy (most of the material in these books has been translated from Russian into English with his name attached as author). We understand that he has been distributing books published by you as part of this scheme, using your name as a reference. If you wish to press charges, please contact this office immediately.

You, upon receiving this new information, write back:

Dear Sirs,

Thank you for this information. I have taken it into account and will modify my activities accordingly. Fortunately, I can make the necessary changes calmly and without any serious upsets. Please convey my best regards to Mr. Powers. I hope the weather is pleasant in Lexington.

Note relayed from W. T. Powers, 62381034:
I don't see reorganization going on here, just the same hierarchy operating successfully with new information. Please send $10,000.

More later. Bill P.

Date: Sat Oct 10, 1992 7:24 am PST
Subject: No reorganization

[From Bill Powers (921010.0930)] Greg Williams (921010) --

Since you don't seem to accept the part of HPCT that depends on the concept of reorganization, I have to ask you how a person can learn to control a new percept, one that this person has "never even considered controlling," without it. The "pandemonium" organization of the HPCT model requires a perceptual function, a comparator, and an output function for each controlled variable. I can see how the necessary perceptual function may already have been formed. But what is the mechanism by which the perceptual signal becomes connected to a comparator? What connects the reference signal to the outputs of the appropriate higher-order systems (remember that this person has never even considered controlling this percept before)? How does the error signal become connected to an output function organized so that its action will result in controlling this percept? And how does the output function acquire that organization?

If this new control process results from hearing some information communicated in words, how does this result in a desire to control something associated with the meanings of this information? Why, upon hearing that the jolly and sympathetic fellow who has kindly invested your money for you is a con man, do you wish to do anything to this con man? Will you continue to perceive this person as a jolly and sympathetic fellow? If not, why not? And if not, by what hierarchical process do you alter this perception? Have you learned some algorithm for changing perceptions? Do you have the ability to resolder your own neural connections at will?

Your proposals require a model entirely different from mine. How about spelling it out?

Best, Bill P.

Date: Sat Oct 10, 1992 9:01 am PST
Subject: Dep[en

[From Bill Powers (921010.0945)] Greg Williams (921010) --

>...the control structure itself at time t1 depends in part on the
>environment prior to t1. (1) WHETHER learning/reorganization takes
>place depends in part on the environment (I think we agree on
>this). (2) WHERE learning/reorganization stops (the resulting new
>control structure) depends in part on the environment (you don't
>agree with this). (3) Attempts to alter another's current reference
>signals "directly" (not involving learning/reorganization) is
>likely to produce conflict in the current control structure (I
>think we agree on this).

1. Learning-reorganization starts when something causes a critical error to become large enough. The external cause itself makes no difference; only the effect on a critical variable and the resulting error matters in setting reorganization into motion. Many different environmental situations may be entirely equivalent in their effects on critical variables (the environment has many more degrees of freedom than there are critical variables). So learning-reorganization does depend in part on the environment, but only in a general existential way. There is no one specific state of the environment on which critical variables depend.

2. Learning-reorganization stops when, for any reason, critical error drops to a small enough value. If the critical error was produced by some effect of the environment on the body, it may disappear when that effect is removed. If the reason for the disappearance was the acquisition (through reorganization) of a learned control system that directly or indirectly removes the environmental disturbance, the learned control system will

continue to function in the same way under the same circumstances that formerly resulted in disturbance of a critical variable. The same environment will then become unable to cause that disturbance of the critical variable; the learned control system will prevent the environment from going into the state that disturbs it.

The criterion for stopping reorganization is zero (or small enough) critical error. This criterion is met only when critical variables are at their respective genetically-defined reference levels. There is no particular state of the environment at which this state occurs. There is no particular learned control structure that will result in correction of critical error -- that is, an infinite number of different control structures could have the same effect of preventing critical error. So while the organization of the learned system does "depend on the environment," the dependence is not systematic but only qualitative. It is ambiguous.

3. Attempts to alter another's ACTIONS are unlikely to result in conflict, because actions normally change as a way of counteracting disturbances. Changing actions requires changing lower-level reference signals. These changes are initiated as a way of counteracting a disturbance at a higher level. The associated perceptual signals remain near in value to the changing reference signals, so control is not disrupted at any level. At the disturbed level, neither the reference signals nor the perceptual signals are significantly changed; the lower-level changes have that purpose.

Conflict results when the disturbance is too large or too fast to be resisted, producing errors that bring other control systems into the picture. It can also result when the lower-level corrective actions are driven into forbidden states. This is not as likely because the higher system being disturbed is not likely to be organized to produce those forbidden states, and thus will not use them to correct its errors. Conflict is produced mainly by an external agent that insists on disturbing a controlled variable, and produces as much force as required to disturb it by a significant amount. In short, by another control system that wants the same controlled variable to be in a different state, or wants something equivalent having the same effect.
--- To say that something depends "in part" on something else is not a step toward precision, but away from it. You can truthfully say that human behavior depends "in part" on "the environment." You can also truthfully say that the climate on Earth depends "in part" on galactic supernovas. Everything depends "in part" on everything else within the same event horizon. But this is not the way to arrive at an understanding of nature. By making more and more general statements, one can eventually arrive at statements that are universally true. "Everything affects everything within the same event horizon" is a universally true statement. But it is also trivial and useless.

It is worse than that, because it implies that all less-general statements consistent with the most general one are also true, and this is not at all the case. There is a very large difference between saying that the form of learned control systems "depends in part" on "the environment" and saying that the environment can be configured in a specific way to determine the form of a learned control system. That is simply not possible; there are too many different ways of controlling that have the same effect. There are too many ways of affecting the environment that would serve to correct the same critical error. The only thing that can be pinned down to any degree is the state at which changes in organization will cease, and that is the state in which critical variables match their reference levels. That is the only predictable outcome of reorganization.

To say that A depends on B is to make a clear statement: given B, one can predict A. To say that A depends "in part" on B is also to make a clear statement: given B, one can predict nothing about A. The qualifier "in part" does not just slightly reduce the amount of dependence. It eliminates dependence altogether.

Best, Bill P.

Date: Sat Oct 10, 1992 9:26 am PST
From: Dag Forssell / MCI ID: 474-2580
TO: SERVER (Ems)
MBX: SERVER@BIOME.BIO.NS.CA

GET MARKEN.DOC
PUB/CSG LIST

```
Date:     Sat Oct 10, 1992  9:30 am  PST
From:     BIOME server
          EMS: INTERNET / MCI ID: 376-5414
          MBX: BIOME-server@biome.bio.ns.ca

TO:     * Dag Forssell / MCI ID: 474-2580
Subject:  MARKEN.DOC
```

The file `MARKEN.DOC' which you requested with the command
        get MARKEN.DOC
cannot be found. Please notify sysop@biome.bio.ns.ca.

        Bill Silvert, System Manager

```
Date:     Sat Oct 10, 1992 11:09 am  PST
From:     Dag Forssell / MCI ID: 474-2580

TO:       Bill                                    (Ems)
          EMS: INTERNET / MCI ID: 376-5414
          MBX: sysop@biome.bio.ns.ca
Subject:  CSG documents
Message-Id: 70921010190907/0004742580NA2EM
```

[From Dag Forssell (921010-direct)]

Bill, I am updating the "starter document" at Gary's request and was interested in Marken.doc for demo disk purposes. Your server just told me that marken.doc cannot be found, ask sysop.

For the benefit of the starter document: How can I request a list of what is currently in the csg/doc subdirectory. Please give a complete command.

For my benefit, and also for clarification of the starter document: How do I request the Marken.doc? I tried: get Marken.doc   If it has been removed, please let me know and send it if you can easily do that for free. If my command was incomplete, then the instructions in the starter document are incomplete.

Thanks for your help.

Dag Forssell
23903 Via Flamenco
Valencia, Ca 91355-2808
Phone (805) 254-1195   Fax (805) 254-7956
Internet: 0004742580@MCIMAIL.COM

```
Date:     Sat Oct 10, 1992 11:12 am  PST
From:     BIOME SysOp
          EMS: INTERNET / MCI ID: 376-5414
          MBX: sysop@biome.bio.ns.ca

TO:     * Dag Forssell / MCI ID: 474-2580
Subject:  BIOME SysOp is not available
```

BIOME SysOp will be away until October 14. He will be attending meetings in Germany and elsewhere.

Mail addressed to this account will not be processed until he returns, although there is a slight possibility of remote login from Europe.

Urgent mail about system operations may be sent to pkeizer@biome.bio.ns.ca (Paul Keizer, 426-3843).

Date: Sat Oct 10, 1992 12:03 pm PST
Subject: Re: Why 99%, hierarchical perception

[Martin Taylor 921010 15:40]   (Rick Marken 921009.1230)

This really has to be brief. Sorry.

>What is the "variance of our estimates"
>by the way? What is the "residual variance" here? the variance of the
>the size of the quadrilateral relative to the reference size? In what
>way is the residual variance "constant" in log size or area units.

The residual variance is the 1% or 2% that you didn't account for. Since the feedback loop function is the same for any level of the reference signal, I assume that the error variance within the ECS, largely due to non-infinite gain, is independent of the level of the reference signal, and that it is a major contributor to the residual variance that you (experimenter) see.

>
>Maybe it's my bias but why not have the reference square vary continuously
>in size?

That would be better, because then you could determine the residual variance as a function of size, and perhaps draw some tighter conclusions than my two-point proposal would have allowed.

>>I think the results would be contaminated by other effects, but the main
>>idea should be valid. It's a very simple kind of case, but it is typical
>>of the sort of thing that psychologists want to know.
>
>Yes, and they miss the whole point. The fact that the person is controlling
>a variable and that it took some doing to figure out what it is just goes
>right by. Early astronomers wanted to know how the patterns they saw in
>the night sky were related to one's personality. I know it's rude but,
>frankly, I could care less what conventional psychologists want to know.
>What they want to know is based on their preconceptions about how people
>work; so they think it is important to know how reinforcement affects
>behavior or which behavioral/cognitive capabilities are in the right
>brain and which are in the left, etc etc.

Well, I want to know how people work, even if you don't. It isn't enough for me to accept that behaviour is controlled perception. I want to know where the signals go (functionally), and what happens if you block this of that path, how to deal with people suffering from stroke, why we have focussed attention and what its limitations are, whether we use internal feedback for short-term memory, and all sorts of questions like that.

Of course the details of what ANYONE wants to know are based on what they think is missing in what they already believe. That's a first-level statement from PCT. So what? If all you are interested in is a succession of demon-strations that perception is controlled, then you are unlikely to find much that is interesting to me. I would like to know where and how, for example, perceptual signals are derived from multisensory inputs (why a sound and a sight seem to come from the same object). The residual variance proposal was a very peripheral example of that. I would like to know whether a multiplier operator is likely to be a common phenomenon in the hierarchy, or whether we should expect logarithmic compression to be common. I'd bet on the latter.

>Since their preconceptions
>about how people work are wrong (because they don't understand the
>consequences of the fact that organisms are locked in a negative feedback
>situation with respect to their environment -- cf. the blind men paper)
>it is only by chance that anything psychologists might want to know about is
>anything more than an illusion from the point of view of PCT.

Some day, I'll write a posting about why conventional psychophysics should for the most part be valid within PCT, but this isn't the time.

Martin

Date: Sat Oct 10, 1992 3:06 pm PST
Subject: percept, the basis ofHPCT HPCT

[From Wayne Hershberger]

Bill Powers:
Bill, your essay on the perceptual basis of HPCT (WTP, 921003) is marvelous. Since I am way behind on my E-mail I have had to limit the threads I read and I overlooked this gem originally. Only after reading and commenting on the later posts by Bruce, Martin and yourself entitled "percepts" (see below) did I belatedly discover your 921003 essay. I think I agree with everything you say (I want to read it several more times to be sure), except one. The thing that bothers me is your wording of the second sentence in the following paragraph--which seems to suggest that a perceptual world passes through the receptors into the brain. I would prefer something like this: "It is that the world we experience directly is ALREADY being shaped by the perceptual processes even as we are experiencing it."

> There is, of course, another interpretation of all this, the one underlying HPCT. It is that the world we experience directly has ALREADY passed through the sensory inputs by the time we experience it. It is the direct manifestation of the brain's way of reading the same world that scientific instruments read. It is the brain's way of realizing, of making apparent, the ordering of the universe. Scientific instruments create a different realization of, we presume, the same underlying order (WTP, 921003).

Martin: Thanks for posting (920924) the Georgopoulos reference.

Martin, Bruce, and Bill: Regarding the term, percept.

Bruce, I, for one, use the term percept to refer to the particulars of experience (objective percepts I call objects for short--as does everyone) essentially as you surmise (92106) a psychologist might use the term; that is:

> conceive : concept :: perceive : percept

This being the case, I take great care to distinguish percepts from neural signals, just as you recommend, and I try to encourage others to do likewise (e.g., my post 920915). Hence, from my perspective, your caveats are well taken.

Further, I believe all four of the terms in your analogy are of the utmost importance in this regard. Although the expression - _conceiving perceiving_ appears to be perfectly reasonable, the expression _conceiving a percept_ (or perceiving a concept) does not. The expression _conceiving a percept_ is an oxymoron because a percept conceived is a concept not a percept. Thus, conceptual models of perception merely assert an equivalence between perceptual and conceptual realizations. On this point I trust we are all agreed. However, it is so easy to confuse this relationship (of a putative equivalence between these two types of realization) with a putative relationship between the various parts of our conceptual models of perception, say, between neurons and photons--or between neural currents and light rays. To call the neural currents percepts is to virtually guarantee the confusion. And to call photons constructs, while claiming that neural impulses are percepts absolutely insures it. So, when I conceive of vision as a perceptual process involving HPCT, I continually remind myself that the system is a conceptual realization from end to end. Further, when I try to imagine (derive?) how this conceptual model accounts for the fact that my visual percepts are perceived as external to my body I am reminded hat the HPCT has an environmental part and that this is not necessarily a mere coincidence.

Dag Forssell: Congratulations to you and Christine on your contract; glad to hear (921001) its going well.

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: Sun Oct 11, 1992 4:13 am PST
Subject: Are We Reorganizing Yet? Should We Vote On It?

From Greg Williams (921011)

Well, maybe at least ONE of us is reorganizing... THREE posts in reply to my last post!

>Bill Powers (921010.0900)

>Dear Mr. Williams,

>The Lexington District Attorney's office wishes to inform you that
>William T. Powers, also known as L. Ron Hubbard III., is presently
>being held by Lexington County police on a charges of extortion, theft
>of services, and fraud in connection with his promotions of
>"dianetics," Rosicrucian therapy, and stock manipulation.....

Hmmm. I know how to deal with obvious hoaxes: ignore them. However, Bill must have thought this important to do for SOME reason, if only to distract his own or others' attention away from the main points of my post. Well... (reorganizing now) taking his project seriously... I first should advise him that to improve the literary quality of his satire he might consider studying the art of that grand master of irony, Mark Twain. There's a complete set of Twain for sale at one of the Durango used book stores.

Even more seriously, I thought Bill had argued way back that running for the exit when somebody yells "Fire!" in a theater involves no reorganization, just continued operation of the person's existing control structure, which includes a reference signal for not getting burned.

[My purported (unreorganizing) answer to the hoax revelation:]

>Thank you for this information. I have taken it into account and will
>modify my activities accordingly. Fortunately, I can make the
>necessary changes calmly and without any serious upsets. Please convey
>my best regards to Mr. Powers. I hope the weather is pleasant in
>Lexington.

I think I would be considerably more upset if I didn't treat this as a hoax. But I can be upset without reorganizing, I think. (I certainly would be upset as I tried to escape from a theater fire). The new information has led me to control differently than if there had been no new information, but that different controlling is via the same old reference signals (get lost money back or don't get burned).

Most seriously of all, what's the point of arguing about whether a particular case of new actions following the presentation of new information involves reorganization or not? I think it does in some cases (i.e., Gary's swimming example) and not in other cases.

>Bill Powers (921010.0930)

>Your proposals require a model entirely different from mine. How about >spelling it out?

I'm not so sure my model is entirely different -- maybe not even different at all, though at least how I interpret the ideological upshot does appear to be different.

When I have time, I'll try to spell it out better than in the past. Right now, I have the October HORTIDEAS, the next BROOKLYN BOTANIC GARDEN NEWSLETTER, CLOSED LOOP, and house-building to deal with. A thorough and careful consideration of alternative models of learning/reorganization would require considerable time, which I don't have right now. I hope you will be patient.

>Bill Powers (921010.0945)

>There is no one specific state of the environment on
>which critical variables depend.

Fine.

>There is no particular learned control structure that will result in
>correction of critical error -- that is, an infinite number of different
>control structures could have the same effect of preventing critical error.

Fine.

>So while the organization of the learned system does
>"depend on the environment," the dependence is not systematic but only
>qualitative. It is ambiguous.

Fine, I think. In different terms (am I with you?), there is no one-to-one mapping between pre-reorganization/during-reorganization environmental states and post-reorganization control structures.

>Attempts to alter another's ACTIONS are unlikely to result in
>conflict, because actions normally change as a way of counteracting
>disturbances. Changing actions requires changing lower-level reference
>signals. These changes are initiated as a way of counteracting a
>disturbance at a higher level. The associated perceptual signals
>remain near in value to the changing reference signals, so control is
>not disrupted at any level. At the disturbed level, neither the
>reference signals nor the perceptual signals are significantly
>changed; the lower-level changes have that purpose.

Yes, I agree. And also that there is another way to control your perceptions which depend on another's actions, without resultant conflict in the other's control structure: arrange the environment so the other's actions are NOT disturbed AND so they result in what you want to perceive (as in Pat serving "healthy" food). But we've already agreed on all that.

>Conflict is produced mainly by an external agent that insists
>on disturbing a controlled variable, and produces as much force as
>required to disturb it by a significant amount. In short, by another
>control system that wants the same controlled variable to be in a
>different state, or wants something equivalent having the same effect.

I agree. This is how the slogan "nobody can control you without overwhelming physical force or threat thereof" unpacks. My point in the previous post -- which I haven't reorganized enough to be distracted from remembering -- is that the slogan, while TRUE, doesn't seem so important when you realize that even though your controlled variables cannot be arbitrarily altered in the short-term without almost certainly giving rise to conflict, your actions (or the effects of those actions as seen by another) CAN be altered. And you might never know what hit you until it is too late. Finally, the other half of the point is that control of your actions by another (speaking loosely) is often something you desire, at least given a background of exchange relations and, yes, in some cases, threats of force.

>To say that something depends "in part" on something else is not a step
>toward precision, but away from it.

To the contrary, it is the first step toward precision, supposing that before saying it, one believed that the something didn't depend AT ALL on the something else. Once it is realized that there is indeed a dependency where before none was thought to exist, it is possible to explore the nature of that dependency.

>"Everything affects everything within the same event
>horizon" is a universally true statement. But it is also trivial and useless.

This is NOT the direction in which I am headed. I am headed toward narrowing the investigation of dependency relations, not broadening the investigation.

>It is worse than that, because it implies that all less-general
>statements consistent with the most general one are also true, and
>this is not at all the case. There is a very large difference between
>saying that the form of learned control systems "depends in part" on
>"the environment" and saying that the environment can be configured in
>a specific way to determine the form of a learned control system. That
>is simply not possible; there are too many different ways of
>controlling that have the same effect. There are too many ways of
>affecting the environment that would serve to correct the same
>critical error. The only thing that can be pinned down to any degree
>is the state at which changes in organization will cease, and that is
>the state in which critical variables match their reference levels.
>That is the only predictable outcome of reorganization.

I specifically dispute that it is impossible for environmental configurations to largely -- but not strictly deterministically -- (one might say "functionally," in the sense that the post-reorganization actions will meet certain criteria so that they "function" as predicted). In Gary's swimming example, the teacher predicts that actions which function to keep the student afloat in deep water will result from reorganization, rather than actions which function, say, to put a part in the student's hair; if the student doesn't drown or give up, the teacher's predictions are correct: the student learned to "swim" rather than to "comb hair."

>To say that A depends on B is to make a clear statement: given B, one
>can predict A.

I don't think so (should we put it to a vote, or just ask a professional philosopher?). To say that A depends on B means that if B changes, A changes, ceteris paribus.

>To say that A depends "in part" on B is also to make a clear statement: given
>B, one can predict nothing about A.

I don't think so (should we put it to a vote, or just ask a professional philosopher?). To say that A depends "in part" on B means that there is also a C such that, if B changes and C doesn't change, A changes, and if C changes and B doesn't change, A changes, ceteris paribus.

>The qualifier "in part" does not just slightly reduce the amount of
>dependence. It eliminates dependence altogether.

Sez Bill Powers, a non-statistical phenomenon -- perhaps even a minority of one?

>Martin Taylor 921010 15:40

>Well, I want to know how people work, even if you don't. It isn't enough for
>me to accept that behaviour is controlled perception. I want to know where
>the signals go (functionally), and what happens if you block this of that
>path, how to deal with people suffering from stroke, why we have focussed
>attention and what its limitations are, whether we use internal feedback for
>short-term memory, and all sorts of questions like that.

>Of course the details of what ANYONE wants to know are based on what they
>think is missing in what they already believe. That's a first-level
>statement from PCT. So what? If all you are interested in is a succession
>of demonstrations that perception is controlled, then you are unlikely to
>find much that is interesting to me. I would like to know where and how,
>for example, perceptual signals are derived from multisensory inputs (why
>a sound and a sight seem to come from the same object).

Well said and worth repeating. Sadly, I predict no reorganization resulting from these comments. I hope my prediction is wrong.

Best wishes,   Greg

Date: Sun Oct 11, 1992 1:50 pm PST

Subject: conceiving percepts; ignoring hoaxes

[From Bill Powers (921011.1500)]

I wrote a lot of longwinded responses to various stuff and decided that it's a waste of time. Brevity rules.

Wayne Hershberger (921010) --

>I would prefer something like this: "It is that the world we experience
>directly is ALREADY being shaped by the perceptual processes even as we are
>experiencing it."

What are these processes and how do they do this shaping?

>Thus, conceptual models of perception merely assert an equivalence between
>perceptual and conceptual realizations. On this point I trust we are all
>agreed.

I would agree if I knew what you meant. What do you mean by "perceptual" and what do you mean by "conceptual?" Pretend I really don't know; you will be right. Are concepts not perceivable?

>The expression _conceiving a percept_ is an oxymoron because
>a percept conceived is a concept not a percept.

Does this suggest a hierarchical relationship? That is, perceiving processes produce percepts, from which a conceiving process can produce a concept?

Greg Williams (921011) --

>Hmmm. I know how to deal with obvious hoaxes: ignore them.

If you really received such a letter and actually believed it, would you still say "Thank you for the information" and simply alter your behavior? Or do you suppose there there might be a tad of affect attached to the situation?

>However, Bill must have thought this important to do for SOME
>reason, if only to distract his own or others' attention away from
>the main points of my post.

I was trying to show you by demonstration (which you foiled by dismissing the letter as a hoax instead of asking what would happen if you believed it) that your example was too intellectualized to be realistic. If you make up examples that have no relationship to real behavior I am going to call you on them.

>>Your proposals require a model entirely different from mine. How
>>about spelling it out?

>I'm not so sure my model is entirely different -- maybe not even
>different at all, though at least how I interpret the ideological
>upshot does appear to be different.

What starts and stops reorganization if critical error doesn't drive it? Why would reorganization take place in response to "Fire" if the person already had the goal of exiting upon hearing "Fire" and immediately did so? How can you be upset (experiencing critical error) without reorganizing? How can you control differently without acquiring a new control system or modifying an old one? How can new information lead to controlling something new without reorganization? How can a person use the fact that another person is reorganizing to control for a particular behavior by the other person? How can a person want another person to control his actions and at the same time want the consequence that those actions are already controlling? Or are you saying that there are actions which are not aimed at controlling anything?

If A is partly determined by B, then to predict A from knowing B you must also know the state of C and all other influences on A, known and unknown, present and future. How is your concept of "ceteris paribus" any different from the failed methods of behavioral science?

>>To say that A depends "in part" on B is also to make a clear
>>statement: given B, one can predict nothing about A.

>I don't think so (should we put it to a vote, or just ask a >professional
philosopher?).

How about reasoning it out? If A depends not only on B but on other variables as well, and you do not know the states of all the other variables on which A depends and cannot predict future states of all those other variables, how can you predict anything about A? How can you determine that an apparent relationship is real? Assuming that you can get data anyway, and make predictions anyway, is the rock on which the ship of psychology has foundered. Why should this method work any better for us?

Martin Taylor (921010.1540) --

>Well, I want to know how people work, even if you don't. It isn't enough
>for me to accept that behaviour is controlled perception. I want to know
>where the signals go (functionally), and what happens if you block this of
>that path, how to deal with people suffering from stroke, why we have
>focussed attention and what its limitations are, whether we use internal
>feedback for short-term memory, and all sorts of questions like that.

Wanting this isn't sufficient to make it possible. We have 50 years of groundwork to lay before any believable answers to such questions can be found. Before you can ask where signals go and what happens if you block this or that path, you have to have a model that is correct. Not just plausible, correct. Throwing together a bunch of suppositions and then using them to make deductions is a total waste of time. Do you want to know how people work, or do you just want to SEEM to know how people work?

Of your statements, Greg said

>Well said and worth repeating. Sadly, I predict no reorganization
>resulting from these comments. I hope my prediction is wrong.

If you get sad ENOUGH, Greg, reorganization will result.

Best to all, Bill P.

Date: Mon Oct 12, 1992 1:11 am PST
Subject: Interaction without theory

[From: Oded Maler 921012] [From Bill Powers (921009.0800)]

* ---
* Control theory fits the world as it is, and people the way they are.
* But an understanding of control theory leads to assigning causality
* differently, and as a result will change the way people do things and
* the reasons for which they do them.

Humans (and surely animals) managed to achieve their goals in the physical world long before Newton. They could interactact with gravity and other forces (and as you would say, reorganize and let their perceptions meet their references) without having a theory for those complex phenomena. Of course, after having the theory they could build tools for influencing the world to a much greater extent.

Now moving to interactions with the living and later the mindfull world: Some people (shepards, politicians, dictators, salesmen, teachers) may succeed in causing individuals and crowds to do certain things because the complex control loop of the influencer is such that it somehow fits some invariant properties of those who are influenced, in the same way that the sensory-motor control loops fits the physical world. It need not be done conciously, nor with a theory of how people are organized. It is "knowing to" and not "knowing why". You may claim that having a theory like PCT will enable more complex interactions in the same way that physical sciences led us to progress from walking to flying.

--Oded

Date: Mon Oct 12, 1992 3:57 am PST
Subject: Reorganization, dependence

From Greg Williams (921012)   >Bill Powers (921011.1500)

>If you really received such a letter and actually believed it, would you
>still say "Thank you for the information" and simply alter your behavior? Or
>do you suppose there there might be a tad of affect attached to the
>situation?

As I said in my previous post, I would probably be upset, but that doesn't necessarily mean I'd be reorganizing. I suspect that affect (of whatever kind) isn't always a sign of reorganization. If you disagree, I guess that's another aspect of HPCT needing to be fleshed out with empirical data.

>What starts and stops reorganization if critical error doesn't drive it?

For the sake of argument, I'll grant that reorganization starts when critical errors get too big (relative to criteria of the organism at that time) and stops when critical errors are reduced sufficiently (relative to criteria of the organism at that time). Where do the criteria come from? They could be inherited and unchanging, or they could been influenced also by the history of the organism itself and/or environmental disturbances. Aside from that second possibility, critical error involves two terms: a reference value for some critical perception and (via subtraction in the PCT model) the perceptual input currently extant; the latter explicitly depends on the independent environment if control is not good (error not equal to zero), as is the case when reorganization is occurring. So, except for solipsists running solely on the "imagination connection," starting and stopping of reorganization depend on BOTH internal AND external events.

>Why would reorganization take place in response to "Fire" if the person
>already had the goal of exiting upon hearing "Fire" and immediately did so?

It wouldn't, I think.

>How can you be upset (experiencing critical error) without reorganizing?

I don't equate being upset and experiencing critical error. Given that: For example, by not being sure you'll make it out of the theater while immediately TRYING TO exit upon hearing "Fire," if you already had the goal of exiting upon hearing "Fire" (and not needing to reorganize). For example, by not being sure that you could control SUCCESSFULLY for getting your money back when being confronted with the revelation of a con scam (which revelation you then believe to be true) while beginning to try to get your money back, if you already had the goal of getting money back which has been taken unlawfully from you (and not needing to reorganize). I don't see reorganization as necessary if the SPECIFIC disturbance to which I suddenly must attend ("offset" to maintain control) doesn't seem intractable, relative to my current control structure. I do believe there can be anxiety WITHOUT reorganization, simply because the success or failure of control isn't always immediate or guaranteed, and it can be difficult to predict the outcome in advance.

>How can you control differently without acquiring a new control system or
>modifying an old one?

"Differently" can be defined in terms of changes in the control structure. You can only control for SOME perceptions with a particular control structure. To control for OTHER perceptions ("controlling differently"), I agree that you MUST "acquire a new control system or modify an old one."

>How can new information lead to controlling something new without
>reorganization?

That depends on what you mean by "controlling something new." If you end up "controlling differently" as I just discussed above, then there must be reorganization. But if the new information results in perceptions which are able to be controlled by the current control structure, but which weren't "actively" being controlled before the new information was perceived, then there need be no reorganization. If I already can control for getting my money back from a crook, when I learn that you're a crook, then I don't need to reorganize, just "actively" (try to) control for getting my money back from YOU, as a PARTICULAR crook -- no reorganization needed. I certainly might be upset because of the (perceived) difficulties of controlling SUCCESSFULLY for getting my money back. What a hassle!

>How can a person use the fact that another person is reorganizing to control
>for a particular behavior by the other person?

Ah, there's the rub: "particular." What I think A can do in some (make that many) cases is arrange B's environment (disturb B) in ways so that B reorganizes and so that the outcome of B's reorganization results in actions by B which are in a class of actions as perceived by A which result in perceptions A is controlling for. If B is ALREADY reorganizing (not due -- in part -- to A's disturbances), I think that at least sometimes A's disturbances during B's reorganization can result in B's post-reorganization actions resulting in perceptions A is controlling for.

How this is possible is by A providing disturbances to B UNTIL B's actions are in the class of actions which result in certain perceptions by A. Note that, to the extent that the PATH of reorganization of B is unpredictable by A, A cannot predict HOW LONG it will take to obtain a successful (to A) outcome of reorganization (a good reason for "teaching" via sequentially arranged "small"-distance reorganizations). But because A is CONTROLLING and hence TRYING AGAIN when B reorganizes the "wrong" way (as A sees it), eventually, B will act as A wants (not "exactly," but within a class of actions) -- unless, of course, B doesn't want to play the game at all, in which case the "teacher"/"counselor"/"parent"/"friend" A would be well-advised to wait until B does want to play the game.

>How can a person want another person to control his actions and at the same
>time want the consequence that those actions are already controlling? Or are
>you saying that there are actions which are not aimed at controlling
anything?

I don't understand the first question here. Please expand on it.

>If A is partly determined by B, then to predict A from knowing B you must
>also know the state of C and all other influences on A, known and unknown,
>present and future.

Yes, to predict EXACTLY. How often in your daily life do you need to predict ANYTHING exactly to see what you want? Never. You don't need to do that, because you are a living control system, not a preprogrammed-output system.

>How is your concept of "ceteris paribus" any different from the failed
>methods of behavioral science?

If A controls for perceptions which depend on B's actions, that control can be successfully achieved, in general, by numerous particular actions of B. A need not predict or control for an EXACT output by B, only for perceiving AN action of B which is A member of the class (defined by A) permitting control by A of certain perceptions of A. A doesn't require PRECISE prediction of B's actions to have successful control of preceptions depending on B's actions.

Where Skinner, in particular, failed was in not understanding that making generative models of the underlying processes provides a means of explaining the CONSTRAINTS on such control. But he wasn't interested in such explanations, so TO HIM it wasn't a failure. If you asked him why THIS rat (recently stuffed with Rat Chow) won't do the Skinner-box tricks that THAT rat (starved for a while) will do, he would have said he didn't care -- but that if you gave him both rats for a couple of days, they would both be doing even MORE tricks than the starved one is doing now. That's "prediction and control"! Skinner-nirvana.

>How about reasoning it out? If A depends not only on B but on other
>variables as well, and you do not know the states of all the other
>variables on which A depends and cannot predict future states of all
>those other variables, how can you predict anything about A?

You can't if all those other variables are making A fluctuate chaotically. But if they DON'T (often the case -- and YOU define the "a lot"), then you can predict future states of A ADEQUATELY FOR YOUR PURPOSES. Trust me, I was trained as a mechanical engineer. How many bridges have you been over which DIDN'T collapse under you? Their designers cannot predict exactly when they WILL collapse (though they are pretty confident they will collapse before t = infinity). And they are also pretty confident that they won't collapse soon, as they define "soon." Maybe they're just fooling themselves. And maybe you're just a minority of one on this issue. Nevertheless, as Steve Earle says (sort of a secular Pascal's wager), "Just because you ain't paranoid don't mean they ain't out to get you."

>How can you determine that an apparent relationship is real?

With control, you don't need to know the causes of disturbances. You only need the connection between your actions and your perceptions to be somewhat non- chaotic. For reorganization, the setting of a problem with a particular CLASS of solutions (i.e., ANY WAY you press that bar gets you food) is what makes it sufficiently non-chaotic.

Best wishes,   Greg

Date: Mon Oct 12, 1992 8:43 am PST
Subject: control theory & behavior; control of behavior

[From Bill Powers (921012.0830)]   Oded Maler (921012) --

>You may claim that having a theory like PCT will enable more
>complex interactions in the same way that physical sciences led us
>to progress from walking to flying.

To be sure, behavior works without benefit of theory, and so do interactions. But in a mindful organism the world is controlled to fit cognitive theories. Theories are part of behavior, not just "about" behavior. We perceive the world through theories, and our actions are often aimed at making the world conform to such theories.

For example, suppose that you have developed a theory to the effect that people with dark skins have lower intelligence than people with white skins. This is a high-level perceptual function: when you look at a person with a dark skin, you will perceive an unintelligent person. That lack of intelligence will be part of what you perceive when you look at that person. It will not seem like a deduction or an opinion: it will be an aspect of that person.

Given that perception, it will soon turn into a reference signal. A black person who seems to act intelligently will not match your experience of such persons, and there will be an error. The next step, of course, is to do something that will correct the error -- modify the perception of the person until the expected impression of intelligence is perceived. So lower-level perceptions will be selected to alter the perception until it is "right" again.

Theories of human nature are intimately connected with the way we deal with other people. If you believe that other people are simply stimulus-response mechanisms, you will try to control their actions in the same way you control anything else. If there are difficulties in achieving control, it will not occur to you that the other person might want to do something other than what you want. You will treat this as a technical problem and look for more effective means of producing the desired result, the way you would do if you wanted to open a door but it was stuck shut. If four hours of food deprivation don't get the result, you try eight, then 24. If noxious stimuli seem required in order to elicit the behavior you want, you use them without compunction. Compunction implies some awareness of the other organism's desires, likes, and dislikes. Such things do not exist in an S-R world. And that is the world that an S-R theorist perceives and controls.

So control theory doesn't just make more "complex" interactions possible. It creates a different perceptual world for the observer and actor, in which the behavior of other people is seen as aimed at satisfying their goals, not as a reaction to passing stimuli. In the behavior of other people you see intentions much like your own. When they resist or push back, you see not a "reaction" but a purposive action, done for a reason, done specifically to prevent your control of something that matters to them. Moreover, you predict that insistence on your own control will not bring you closer to getting what you

want, but will elicit an escalation of the other's resistance, possibly turning it into active attack. From the CT standpoint, the violence that goes on between people is simply the natural result of what they're trying to do to each other. It isn't caused by aggressive impulses, territoriality, base motives, or stupidity. It's caused by trying to control each other. If they stopped trying to control each other, the violence would stop. If they had a better theory of human nature, they would understand what's causing the problem.

Greg Williams (921012) --

I suppose that there is a scale of upsetness, and that at the lower amounts there's simply a mild error that can be tolerated or corrected at leisure. I think that reorganization operates on the basis of duration times intensity of error; that's just saying that the reorganizing system perceives on a longer time-scale than the hierarchy does. A significant chronic error in a critical variable is required to start reorganization. I think such errors occur fairly often, and that they are associated with internal states that we experience as emotions (not necessarily with negative labels). But there's a lot of room for quantitative disagreement here, and no way to settle it but getting data.

As to criteria for reorganization, I think the basic ones have to be built in along with their reference levels. This doesn't rule out others. The assumption of built-in-ness, however, is based on evolutionary and co-evolutionary grounds. The reorganizing system has to be operational before any control systems become organized, and before any perceptions higher than intensities exist. Furthermore, it has to be able to produce competent control systems in any environment that might be experienced. Evolution can't anticipate the details, only whatever is consistent over hundreds of thousands of years.

I think we have different notions of what constitutes critical error, or perhaps we simply have different experiences. If I try to do something important to me, and anticipate that it won't work, I experience some conflict -- hestitation -- about trying it at all. I feel emotions. I'm prepared for action but am holding myself back. I think this state amounts to a little bit of critical error. Maybe not a big important one, but enough to make me start thinking of some alternative, casting around for a way that feels better. And I would count that as a little bit of reorganization.

>If I already can control for getting my money back from a crook,
>when I learn that you're a crook, then I don't need to reorganize,
>just "actively" (try to) control for getting my money back from >YOU,
as a PARTICULAR crook -- no reorganization needed.

But how about the reorganization needed to perceive me as a crook instead of how you were perceiving me before you got the new information? I don't think that such a reorganization could occur unless there was some serious kind of upset to motivate it. Of course if you don't think it's reasonable to equate shock, dismay, a feeling of being betrayed, and anger with a disturbance of critical variables, then I guess you would believe that you could make this change without reorganizing.

When I ask you for your telephone number and you give it to me, I can do a behavior that I couldn't do before: call you up. I think that only memory is involved here, no reorganization. So I don't dispute that control can be "facilitated" by getting new information which is handled by existing control systems. But I don't think that this sort of facilitation has any deep theoretical signficance. And I don't think you can reduce situations like being told you're in danger from a fire to the same situation as being told a telephone number.

>>How can a person use the fact that another person is reorganizing
>>to control for a particular behavior by the other person?

>Ah, there's the rub: "particular."

My point exactly.

>What I think A can do in some (make that many) cases is arrange B's
>environment (disturb B) in ways so that B reorganizes and so that

>the outcome of B's reorganization results in actions by B which are
>in a class of actions as perceived by A which result in perceptions
>A is controlling for.

I wish you wouldn't use "controlling for" in this loose way, when what you mean is "wishes to see." You can control for something only when your actions have a systematic effect on it and maintain it near your reference level. B can arrange A's environment in a way that B thinks will have some chance of producing a behavior that B wants to see. If A produces that behavior, B will be gratified, but will not have any control; A could do something else, and B would have no way of altering that. The best that B could do would be to predict that over many occasions and with many A's, arranging the environment in a particular way will produce some percentage of outcomes of reorganization that will fit B's desires. To get any better results than this, B would have to have extensive control over A and A's environment, as in a Skinner box. There's just no innocuous way to accomplish what you're describing.

>If B is ALREADY reorganizing (not due -- in part -- to A's
>disturbances), I think that at least sometimes A's disturbances
>during B's reorganization can result in B's post-reorganization
>actions resulting in perceptions A is controlling for.

But I thought we were talking about control, purposeful influence, not statistical effects. If you loosen the concept of control to include poor, chancy, unreliable control, in which you can never be sure what the effect of your action really was, then everything becomes possible. The con man can be sure of fooling the mark if he can try his pitch on as many people as he likes and count only the successes. The advertiser can claim to control buying behavior if out of 20,000,000 people who see the ad, 2000 of them buy the car (and only 1500 of them would have bought it anyway).

>>How can a person want another person to control his actions and at
>>the same time want the consequence that those actions are already
>>controlling?

>I don't understand the first question here. Please expand on it.

You referred to WANTING another person to control your action. Actions are produced only to control something other than the action. If you want someone else to control your action, this means that you have a preferred state for your action, at which you want the other to control it. But that action has to be freely variable in order to combat unpredictable disturbances; you can't have a preferred state for an action at the same time you're using it to control something else. It doesn't matter whether you want to control the action yourself or to have someone else control it; controlling the action will destroy your ability to control the variable it was being used to control.

The only way the other person can control your action is to disturb the controlled variable, acting in parallel with your action. This can't improve your ability to control the variable. Whether it aids your effort or opposes it, your effort will change so as to oppose the other person's contribution to the controlled variable's state. If the other's effort is aiding, you will relax. You couldn't help relaxing unless you reorganized your control system. Disturbances are always opposed, whether they're meant to be helpful or not. Controlling your action in this way depends totally on the degree of control you already have over the variable being controlled. There's no way another person can help you control a variable by pushing on it -- not without eliciting opposition.

Of course that opposition might be OK with the helper. The helper might not mind if you relax and leave some of the load for the helper to support. Of course then what the helper is teaching the helpee is not how to control better, but how to control worse.

Hook one end of the rubber-bands over a stationary object, then pull the knot to a target position. Have someone else, with a third rubber band attached to the knot, help you pull. See what happens to your end of your rubber band.

>>If A is partly determined by B, then to predict A from knowing B
>>you must also know the state of C and all other influences on A,
>>known and unknown, present and future.

>Yes, to predict EXACTLY. How often in your daily life do you need >to
predict ANYTHING exactly to see what you want? Never.

I agree that we (almost) never need to predict exactly, or at all, in order to see what we want. We just control, which doesn't require any prediction. But when you do predict, I think you want to get closer than you could do by estimating A from B when all the other variables that affect A are unpredictable.

I'm objecting to your method because it's basically the same method used in standard psychology. The errors of prediction achievable through sophisticated application of advanced statistical methods, using many trials and many subjects, are in the hundreds of percent, if not thousands, in any specific instance. If you think that kind of predictability is good enough for ordinary behavior, why would we need control systems? I freely admit that people DO use this sort of prediction, and that they DO think it's good enough. That's a delusion, but a popular one.

>You can't [predict] if all those other variables are making A
>fluctuate chaotically. But if they DON'T (often the case -- and YOU
>define the "a lot"), then you can predict future states of A
>ADEQUATELY FOR YOUR PURPOSES. Trust me, I was trained as a
>mechanical engineer. How many bridges have you been over which
>DIDN'T collapse under you?

If we're talking about the inanimate environment, prediction works fine. Trivial internal structure, no goals, no levels, no reorganization. Science hasn't had much trouble with prediction in that context. I trust you as a mechanical engineer. But if you're going to apply the same approach to people, I think you're in for a disapppointment.

>>How can you determine that an apparent relationship is real?

>With control, you don't need to know the causes of disturbances. >You only need the connection between your actions and your >perceptions to be somewhat non-chaotic.

True.

>For reorganization, the setting of a problem with a particular >CLASS of solutions (i.e., ANY WAY you press that bar gets you food) >is what makes it sufficiently non-chaotic.

But you have to make sure that the rat doesn't escape and has no other source of food and is hungry. You can get a reorganizing system to solve YOUR problem only when you already have control over the organism in most other important ways. You can't just walk up to a stranger and set a problem and expect it to be solved: "You are one of two prisoners accused of a crime. If you confess and the other doesn't, you will get a stiff sentence..."

Most of the methods you propose for controlling other people, or even predicting their behavior, simply won't work in the wild. Most of them depend on establishing background conditions that could reliably be established only by brute force: solve this problem or I will shoot you. I believe that such methods of control do exist and are applied successfully. But no method of control applied to a subject with, as it were, a gun to the head has much theoretical or practical signficance when the gun is removed. The gun makes all techniques work. Try some examples in which there is no gun, explicit or implicit, and you will see the true locus of control.

Best to all,   Bill P.

Date: Mon Oct 12, 1992 10:49 am PST
Subject: Rush to knowledge

[From Rick Marken (921012.1130)]   Martin Taylor (921010 15:40) --

I still have absolutely no idea what you expect to find by analyzing "residual variance" or why you expect to find it. How about actually writing some equations or, better, some model code so that I can understand how your proposal relates to a functional control model. I still believe that the statistical analysis you are so earnestly pursuing is worthless; but I am willing to be convinced otherwise, if you want to try. But I certainly don't mind if you want to have a good time doing these kinds of analyses --

they are familiar and comfortable for you. But sometimes, in order to understand new things, you have to try what is not familiar and comfortable. It's called reorganization. It happens when current control systems are not working. Your "data analysis" systems seems to be working just fine so, mazel tov.

Martin also said to me:

>Well, I want to know how people work, even if you don't. It isn't enough for
>me to accept that behaviour is controlled perception. I want to know where
>the signals go (functionally), and what happens if you block this of that
>path, how to deal with people suffering from stroke, why we have focussed
>attention and what its limitations are, whether we use internal feedback
>for short-term memory, and all sorts of questions like that.

In support of Martin's comments, Greg said:

>Well said and worth repeating. Sadly, I predict no reorganization
>resulting from these comments. I hope my prediction is wrong.

Bill Powers (921011.1500) replied to Martin (and Greg):

>Wanting this isn't sufficient to make it possible. We have 50 years of
>groundwork to lay before any believable answers to such questions can be
>found.

>Before you can ask where signals go and what happens if you block this
>or that path, you have to have a model that is correct. Not just plausible,
>correct. Throwing together a bunch of suppositions and then using them to
>make deductions is a total waste of time. Do you want to know how people
>work, or do you just want to SEEM to know how people work?

To which I can only add an enthusiastic "here, here". PCT can't be much fun for people who already know what they want to find out about how people work. It's no fun because people don't work that way.

Have a nice week   Rick

Date: Mon Oct 12, 1992 11:12 am PST
Subject: Re: conceiving percepts; ignoring hoaxes

[Martin Taylor 921012 15:00]   (Bill Powers 921011.1500)

> Before you can ask where signals go and what happens if you block this or
>that path, you have to have a model that is correct. Not just plausible,
>correct. ... Do you want to know how people work, or do you just want to
>SEEM to know how people work?

Not being God, I'll settle for the latter. It's all a human can aspire to.

How does a human START with a model that is correct, rather than plausible? I'll go with the experimental method, thanks. Start with plausible models and see which accord better with the data. But I know I'll never have a model that is both correct and comprehensible, let alone one that I KNOW to be correct.

I do hope that your reorganization processes do at some time come to alter your thinking about prediction and information. I take it that you have no insurance?

Happy Thanksgiving, Canada.   Martin

Date: Mon Oct 12, 1992 11:26 am PST
Subject: Re: Reorganization, dependence

[Martin Taylor 921012 15:05]   (Greg Williams 921012)

Good comments to Bill on partial prediction. A technical quibble, though. You equate "chaotic" with what I would call "random." It is because the world IS chaotic, not because it is not, that the engineer can predict that the bridge will not fall "soon" but will fall at some indeterminate time in the future.

Chaotic means, loosely, that short term prediction is normally pretty good, but long-term prediction is impossible. My weather forecast for the next millisecond is essentially perfect, for the next hour not bad, for the next day better than chance (I can often decide reasonably that we will have a picnic tomorrow), and for next week pretty useless. Professional forecasters do better, but can't go beyond about a week. The weather is a chaotic system.

In a random system, knowing the state at moment t1 gives you no information about the state at any later moment t2, no matter how close t2 is to t1.

It is the chaotic nature of the world that permits PCT to work, and at the same time allows short-term planning and habitual action patterns to be useful. If the world were not chaotic, PCT would be unnecessary, since pure planning could work. If it were random, PCT could not work. PCT depends on the chaos of the world for its ability to stabilize perceptions.

Martin

Date: Mon Oct 12, 1992 12:44 pm PST
Subject: Nonlinear models; modeling

[From Bill Powers (921012.1400)] Rick Marken (921012.1130) --

Rick, you're giving Martin a pretty hard time for (I think) the wrong reason.

When you get the unexplained variance as low as it will go using a linear model, the next thing to try is a nonlinear one. To test a nonlinear model you have to use more than one value of the reference signal, because if the perception is nonlinear, the loop gain will vary with signal amplitude. Martin suggested trying logarithmic perceptual functions, which isn't a bad idea, considering Weber and Fechner, and Stevens.

Of course when you get to the point where the residual errors of prediction are only a few percent, trying to reduce them further can easily get you into what Runkel calls "fine slicing." Bringing in nonlinearities introduces more adjustable parameters (a*log(bx)) instead of just bx), which you have to try to estimate from the little bit of variance that's left, which can make ALL the parameters less accurately known, and so on to meaningless detail.

I'm satisfied when we get the variance into what I consider a "normal" range of experimental error; it doesn't seem worthwhile extending little pseudopods of precision into a forest of imprecision. But when I sweep a floor I do a raster scan once, while Mary likes to do a bit here and a bit there. The floor gets swept in either case.

The only serious problem I see with spending a lot of time on one narrow part of the organization is that when you get the other parts into better shape, you can easily find that you were pursuing a wrong idea, so all that effort will be wasted. What if we decide to use the "new" model in which lower order error signals instead of perceptual signals go to the higher levels?

Martin Taylor (921012.1500) --

>How does a human START with a model that is correct, rather than
>plausible?

I meant, before we start elaborating on a model, we should be sure that what we're elaborating on is a solid foundation. But see the above post to Rick; this is partly a matter of taste. It's my personal conviction that running ahead of the data too far is usually a waste of time, but some people are better guessers than others, so ...

I agree: we should start with plausible models and see which accord better with the data. That's the beginning. Then we should ask why the best of the models doesn't work any better than it does, and try to refine it and make it predict new data as closely as possible, up to a cutoff point where further refinements yield diminishing returns.

>I do hope that your reorganization processes do at some time come
>to alter your thinking about prediction and information. I take it
>that you have no insurance?

I waste my money on insurance like everyone else. Can't say that I do much calculation of probabilities when I buy it. I do preventive maintenance like oil changes, though I don't know how often I really should do them (probably not as often as I do). I subscribe to magazines which I expect to show up on time. I lay in wood for the winter. Let's see, what else? Oh, I drive on highways, assuming other drivers will stay in their lanes (that is, I ignore them -- not really a prediction). I vote for people about whom I know almost nothing. There's not really much more that I actually do any predicting about. Most of the time I'm just trying to satisfy my goals and whims and cope with the variance of the environment as it comes along. Control is very handy, in that you don't have to know what's going to happen next; you can handle most disturbances when they come up.

Happy Thanksgiving, Canadians.

Best to all, Bill P.

Date: Mon Oct 12, 1992 2:01 pm PST
Subject: Review comments

Hello, Martin,

I think you should reread the part of your review in which you referred to the Nyquist criterion (with which I am quite familiar). We reported that we sampled the data 30 times per second. You said that this was impossible. In other words, you said we could not have sampled the data 30 times per second. Therefore we must have been lying. That is what we objected to. Our experiment iterated 30 times per second, sampling handle position and recomputing cursor position on every iteration. This was considerably faster than the Nyquist criterion would suggest for the traditional 2.5 Hz bandwidth of manual tracking, but was necessary to show enough detail to detect all errors between the real person and the model's behavior. A nominal bandwidth of 2.5 Hz in tracking performance does not include all rapid variations that are actually produced by human subjects.

You seem to have decided sometime lately that S-R experiments are just control experiments with the loop broken. This is not at all the relationship I see between an S-R experiment and a parallel control experiment.

S-R psychology arose from the observation that when certain events impinge on organisms, the organisms seem to respond to the events. In an intact control system, with the loop closed, there is an exactly parallel phenomenon: when a disturbance occurs, its effects on a controlled variable result in a change in action that cancels the effects.

So far you seem to be thinking of a disturbance as a controlled variable that is arbitrarily altered. You asked me, in fact, what difference loose or tight coupling of the disturbing variable to the controlled variable made; this shows me that you don't yet get the point. Loose coupling is required in order to allow the disturbing variable to change by an easily-observable amount, while the control system prevents the controlled variable from changing significantly. The disturbing variable is a physical variable different from the controlled variable. It is linked to the controlled variable in some way that has "give" in it. Thus for a very high-gain control system, it is possible for the disturbing variable to change over a wide range while the controlled variable's changes are too small to measure. At the same time, the output will change over a range comparable to that of the disturbing variable, opposing its effects. If one actually seized the controlled variable and tried to make it change by any significant amount, the output of the system would immediately slam up against the stops. Even in an ordinary behavioral control system with a gain on the order of 30, very small changes in the controlled variable (3 percent of the maximum value of the perceptual signal) result in full-scale changes in the action that controls it. Actually breaking the control loop would give you a system that is hypersensitive to the smallest changes in "stimulation." This is not what is observed in S-R experiments.

I believe that the correct interpretation of most S-R experiments is that they are manipulating disturbing variables that are only loosely linked to the actual controlled variables. The controlled variables, because they are kept by the "response" from changing, are discarded by the standard statistical analysis; they do not correlate with either the action or the disturbance. The disturbance and the response do correlate because the response is opposing the effect of the disturbance on the controlled variable. Of course, not knowing that there is a controlled variable, the S-R experimenter does not pick a measure of either the disturbance or the response that is appropriate to the common effect on the controlled variable; the correlations actually found are far lower than they would be if the correct measures were used. This loss of precision is exacerbated by the fact that not knowing of the controlled variable, the experimenter can't protect that variable from other disturbances that will also produce opposing changes in action.

So standard "scientific method" is perfectly designed to discard controlled variables, and to select as independent variables remote disturbances that have effects on controlled variables -- or would have effects if it were not for the response. S-R experiments actually explore the relationship between disturbances that are loosely coupled to controlled variables and the actions that prevent the controlled variables from being disturbed. The disturbing variable is interpreted as a "stimulus," and it is assumed that the organism is sensing it directly. So it appears that a stimulus is causing a response. This is what I have referred to as the "behavioral illusion."

In my conversation with you I said that our target audience consisted of psychologists who were sufficiently dissatisfied with their science to be looking for a new approach. I did not say that it consisted of people "naive with respect to psychology but who might happen to read that journal." You are distorting my words to make your position look more reasonable.

Bill.

Date: Mon Oct 12, 1992 2:05 pm PST
Subject: Nonlinear models; modeling

[From Rick Marken (921012.0300)] Bill Powers (921012.1400) --

>Rick, you're giving Martin a pretty hard time for (I think) the wrong reason.

Me? Giving a hard time? Wrong reason!!

>Martin suggested trying logarithmic perceptual functions, which isn't a
>bad idea, considering Weber and Fechner, and Stevens.

I was not objecting to the log function idea. I'm giving Martin a hard time because a) it's fun and b) because he seemed to suggest that one could tell from the nature of the residual variance itself something about the nature of the required non-linearities in the perceptual function. My "cheerfully ignorant" impression was that Martin's suggestion was based on the way this is done in conventional statistical research (which assumes y = f(x1,x2,...xn)) and accounts for the residual error variance by adding in the non-linear predictors, x^2, x^3, etc. And you can tell from the residuals what terms might be best to include in the regression. I don't think you can look at the residuals and tell much about how to change the control model; but, as I said, I'd be happy to find out how it can be done. I just didn't follow Martin's description. So I am not objecting to any particular proposals about how to change the model to improve prediciton. I am questioning the ability to go from an analysis of the residuals (based on an input-output model of the system) to an understanding of what it is about the system model that should be changed. I'm not saying it can't be done; I'm just asking how. I am trying to encourage an approach to research based on modelling rather than curve fitting. If Martin can show how an analysis of residuals such as he suggests can be used to make coherent improvements to a closed loop behavioral model, then I think that would be a nice contribution to PCT methodology.

Best regards Rick

Date: Mon Oct 12, 1992 6:40 pm PST
Subject: I saw the "impossible" and lived to tell about it

From Greg Williams (921012 - 2)

Belated thanks to Bruce Nevin (sorry, I lost his post's date-time), who mentioned the hypnotic techniques of Milton Erickson. There is a new book out on Erickson by Bill O'Hanlon which might have a lot to say regarding purposive influence. I've ordered it and might post a review of it in the near future. It SOUNDS less obscure than some of the other Erickson books....

>Bill Powers (921012.0830)

>I suppose that there is a scale of upsetness, and that at the lower amounts
>there's simply a mild error that can be tolerated or corrected at leisure.

And I suppose that there are more than a single scale of upsetnesses -- some upsetnesses occurring with reorganization, some not. One kind of upsetness (which can be more or less in amount) can occur when the success of controlling is in doubt, but reorganization isn't triggered. If you don't think this sort of upsetness is reasonable to postulate, please consider again the example of successfully exiting a theater after "Fire!" has been yelled, without (we both apparently agree) reorganization; it is unlikely, I believe, that the exiter would not have been at least a bit upset during the exiting.

>But there's a lot of room for quantitative disagreement here, and no way to
>settle it but getting data.

Indeed.

>As to criteria for reorganization, I think the basic ones have to be
>built in along with their reference levels. This doesn't rule out
>others. The assumption of built-in-ness, however, is based on
>evolutionary and co-evolutionary grounds. The reorganizing system has
>to be operational before any control systems become organized, and
>before any perceptions higher than intensities exist. Furthermore, it
>has to be able to produce competent control systems in any environment
>that might be experienced. Evolution can't anticipate the details,
>only whatever is consistent over hundreds of thousands of years.

I have no problem with this. I do think, though, that it might be possible for acquired criteria to override the built-in ones. I don't think that possibility is a problem for either of our viewpoints.

>... enough to make me start thinking of some alternative, casting around for
>a way that feels better. And I would count that as a little bit of
>reorganization.

I don't have a strong objection to this. Maybe it would be best to revise the PCT-explanation of the "Fire!"/exiting example to say that reorganization DID take place? That would be OK with me. In other words, I see no objection if practically every time there was any "new" information perceived, there was at least a little reorganization. But, as you said above, data are needed to settle this issue.

>But how about the reorganization needed to perceive me as a crook instead
>of how you were perceiving me before you got the new information?

Or how about reorganizing to perceive that theater as a fire trap? Again, I'll be happy with consistency either way: exiting and con-realizing WITHOUT ANY reorganization, or both WITH SOME (perhaps minimal) reorganization. My problem with a lack of consistency is that you seem to want to treat the two examples as fundmentally different, yet I don't see a fundamental difference.

>When I ask you for your telephone number and you give it to me, I can
>do a behavior that I couldn't do before: call you up. I think that
>only memory is involved here, no reorganization. So I don't dispute
>that control can be "facilitated" by getting new information which is
>handled by existing control systems. But I don't think that this sort
>of facilitation has any deep theoretical signficance. And I don't
>think you can reduce situations like being told you're in danger from
>a fire to the same situation as being told a telephone number.

I've been thinking until right now that you said that the exiting after "Fire!" example was NOT an example of reorganization. Do you now claim or have you been claiming it IS an example of reorganization? If so, please excuse my mistake and accusation of inconsistency.

If memory, not reorganization, is involved in a particular instance of "facilitation" (or, more generally, "purposive influence"), then that instance is a kind of "rubber-banding," which might not have, for you, what you call "deep theoretical significance," but certainly has great practical significance AND scientific significance, in my opinion.

GW>>What I think A can do in some (make that many) cases is arrange B's
GW>>environment (disturb B) in ways so that B reorganizes and so that
GW>>the outcome of B's reorganization results in actions by B which are
GW>>in a class of actions as perceived by A which result in perceptions
GW>>A is controlling for.

>I wish you wouldn't use "controlling for" in this loose way, when what
>you mean is "wishes to see." You can control for something only when
>your actions have a systematic effect on it and maintain it near your
>reference level. B can arrange A's environment in a way that B thinks
>will have some chance of producing a behavior that B wants to see. If
>A produces that behavior, B will be gratified, but will not have any
>control; A could do something else, and B would have no way of
>altering that. The best that B could do would be to predict that over
>many occasions and with many A's, arranging the environment in a
>particular way will produce some percentage of outcomes of
>reorganization that will fit B's desires. To get any better results
>than this, B would have to have extensive control over A and A's
>environment, as in a Skinner box. There's just no innocuous way to
>accomplish what you're describing.

This is the crux of our dispute. I claim that this is truly CONTROLLING FOR, not just "wishes to see." B arranges A's environment so as to encourage a class of actions by A which B wants to see. If A doesn't perform actions in the class defined by B, then B RE-arranges A's environment. And so on, until A does actions in the class defined by B, or B gives up. IN PRACTICE, I see that this works much of the time: A indeed does perform actions B wants to see, and often within a short time. In principle, there is no difference between this sort of control and the control of a cursor subject to a "hidden" disturbance -- in both cases, what is tending to thwart control cannot be "seen." But B can do quite a bit to get around the problem, like ask A, "Are you sure you REALLY want to learn to swim, rather than to comb your hair?" and B can charge A a stiff fee for "teaching" A "swimming." Still, there is always the possibility that A will lie about his/her motives and is paying a stiff fee to get B alone so he/she can drown him/her. There is no difference in principle between those sorts of possibilities and the possibility that the computer in the cursor-control trials will break, so that control is impossible. Such is life. Even non-social life: the gravitational constant might start fluctuating wildly at 2PM today. And a LOT of controlling would suddenly become quite difficult.

However, our models of physics suggest that wild fluctuations in the gravitational constant, beginning at 2 PM today, are unlikely. And -- here is where Skinner feared to tread -- PCT models suggest what constraints are important in determining the likely success or failure of attempts at "purposive influencing." PCT explains why it is easy for an experimenter to control for seeing actions (which are in a functional class of actions the experimenter has defined, like "actions which press the lever in the box, which happens to release Rat Chow") of a starved rat in a Skinner box (Skinner didn't know WHY it is easy). PCT explains why it is not quite so easy for a "teacher" to control the actions (which are in a functional class... called "swimming actions") of a person who comes to a "teacher" and pays money to be "taught" to "swim" (again, Skinner didn't know WHY it isn't quite so easy). PCT does NOT say that "teaching swimming" is impossible IN GENERAL, but it does explain why it can fail in some cases. It can even fail in ways in which control depending only on non-living things cannot, i.e., if the "student" starts controlling for NOT learning to swim; presumably, fluctuations in the gravitational constant wouldn't be purposive.

I suppose "innocuous" is in the eye of the beholder. Exchange relations seem rather innocuous to me, but maybe I'm just not enough of a revolutionary. Most of the time, I don't mind not being able to spend other people's money. But some people do mind that "imposition," much of the time -- I realize that. I'm not a Pollyana: NOT ALL social interactions are "win-win". But I don't think all are "lose-lose" or "win-lose," either.

>The con man can be sure of fooling the mark if he can try
>his pitch on as many people as he likes and count only the successes.

The big-con artists do not operate on a statistical basis. They take time and pains to model the control structure of each potential mark, and give up (as PCT suggests they should) if the mark doesn't appear to want what they need the mark to want, in order for their (the con artists') controlling, which depends on the mark's actions, to be successful. Of course, there ARE controllers who DO make statistical models of control structures at the population level: advertisers, politicians, economists, movie directors, magicians, and others.

BP>>How can a person want another person to control his actions and at
BP>>the same time want the consequence that those actions are already
BP>>controlling?

GW>I don't understand the first question here. Please expand on it.

>You referred to WANTING another person to control your action.

In what context?

>Actions are produced only to control something other than the action. If you
>want someone else to control your action, this means that you have a
>preferred state for your action, at which you want the other to control it.
>But that action has to be freely variable in order to combat unpredictable
>disturbances; you can't have a preferred state for an action at the same time
>you're using it to control something else. It doesn't matter whether you want
>to control the action yourself or to have someone else control it;
>controlling the action will destroy your ability to control the variable it
>was being used to control.

I don't see anything wrong with what you say here. Maybe I meant somebody wanting somebody else to "teach" him/her something new, so the first party ended up reorganized. Or maybe I just got confused. That IS possible. (:->)

GW>>For reorganization, the setting of a problem with a particular
GW>>CLASS of solutions (i.e., ANY WAY you press that bar gets you food)
GW>>is what makes it sufficiently non-chaotic.

>But you have to make sure that the rat doesn't escape and has no other
>source of food and is hungry.

Or you have to make sure that the person wants to learn to "swim" with you as "teacher." (I suggest asking, rather than threatening.) Or you have to make sure that the computer in the cursor-control experiment won't self-destruct.

>You can get a reorganizing system to solve YOUR problem only when you already
>have control over the organism in most other important ways.

Or if your model of the reorganizing system's desire to solve ITS problem is accurate. If the system says (as I once did to a swimming teacher), "No way I'm going to try to swim," you had better say, "Next student, please!"

>You can't just walk up to a stranger and set a problem and expect it to be
>solved...

Exactly. PCT explains why (Skinner couldn't). PCT also explains why you CAN walk up to a NON-stranger and set SOME KINDS of problems -- depending on the non-stranger's control structure (as modeled by you) -- and expect them to be solved.

>Most of the methods you propose for controlling other people, or even
>predicting their behavior, simply won't work in the wild.

I disagree. I see them working "in the wild." (Yes, even AWAY from wild Black Lick Hollow.)

>Most of them depend on establishing background conditions that could reliably
>be established only by brute force: solve this problem or I will shoot you.

I disagree. Many depend on background conditions which could reliably be established by specialization of professions: I'll "help" you solve the problem you want to solve if you pay me. Granted, there's some brute force underlying private-property economics, but it isn't at the same level you're talking about. The point is, the problem being solved by guided reorganization is the REORGANIZER'S problem (whether or not the third graders realize it; some don't until much later (age 26, tax time: "Why didn't I study those multiplication tables?"), which is why there are truant officers, maybe even carrying guns).

Many depend on non-economic reciprocities. Like wanting to feel nice in exchange for lifting a little old lady's suitcase (remember?).

>I believe that such methods of control do exist and are applied successfully.
>But no method of control applied to a subject with, as it were, a gun to the
>head has much theoretical or practical signficance when the gun is removed.
>The gun makes all techniques work. Try some examples in which there is no
>gun, explicit or implicit, and you will see the true locus of control.

Today my son Evan was having a problem with his new birthday present, a radio- controlled truck. He asked me to help him figure out what was wrong with the transmitter. Some experiments guided by me showed a weak battery. Next time he'll be able to cure the malady himself. No, he didn't hold a gun to MY head, either. We BOTH got to where we wanted to be. I saw me controlling for him learning how to solve the problem in the future. I saw him controlling for a solved problem ASAP. In a couple of days (or sooner -- the truck's batteries wear out pretty quickly!), I'll be happy he can solve the problem, and so will he.

>Martin Taylor 921012 15:05

>Good comments to Bill on partial prediction. A technical quibble, though.
>You equate "chaotic" with what I would call "random." It is because the
>world IS chaotic, not because it is not, that the engineer can predict that
>the bridge will not fall "soon" but will fall at some indeterminate time
>in the future.

I see what you mean, and I agree that I should have said "random," or at least "pseudorandom."

Best wishes,   Greg

Date: Mon Oct 12, 1992 8:12 pm PST
Subject: digital dribble

to dennis delprato: thankyou for sending the paper and accepting the interview.

to the Player marken: please send me a copy of the regretably rejected paper.

to w.t. bourbon: please send me a copy of the "to control or to be controlled" paper, since it is fundamental to my sparta-analysis.

to martin: question- does not pct predict events in the future as well as in the present (i.e. that unless some overwhelming disturbance stops me through raw power or violation of system integrity - an ogre or injected lye - then i will sucessfully drive to the store. this is often far into means of course. but it also makes "immediate" predictions if the goal is known and functional characterisics of the system AND (that's a big and) the disturbances that be incountered then pct also makes accurate predictions, this is of course if you don't care whether i drive with my hands or with my feet (as difficult/silly as that may be). i'm probably rambling and could of well missed your point so i'll shut up now.

i.n.kurtzer

Date: Tue Oct 13, 1992 6:29 am PST
Subject: RE: I saw the "impossible" and lived to tell about it

> From Greg Williams (921012 - 2)
>
> This is the crux of our dispute. I claim that this is truly CONTROLLING FOR,
> not just "wishes to see." B arranges A's environment so as to encourage a
> class of actions by A which B wants to see. If A doesn't perform actions in
> the class defined by B, then B RE-arranges A's environment. And so on, until A
> does actions in the class defined by B, or B gives up.

All social transactions (work, love, marriage, et al.) occur in the realm of actions. Control system A needs control system B to perform some useful (to it) action. There are obviously a wide range of purposes (aka motivations or causes) driving control system B to perform these actions.

However, it is our actions that are judged and endure in the world, not our control structure nor purposes. In addition, actions are the visible (tip of the iceberg) portion of our control structure. So the crux of my argument is that it is (perhaps unintentional) actions (byproducts of control) which drive the world.

Curt

****
Curt McNamara (mcnamara@mgi.com) |"the mome rath isn't born that
Mgmt. Graphics, Inc. |could outgrabe me."
1401 E. 79th St. | Nicol Williamson
Mpls., MN 55425 |
****

Date: Tue Oct 13, 1992 10:36 am PST
Subject: Statistics; Loop gain; actions/intentions

[From Bill Powers (921013.0930)] Rick Marken (921012) --

It strikes me that one problem with "residuals" and all that is simply that the wrong model is used (as you say). Is there anything to prevent you from doing statistical manipulations using a closed-loop model instead of an open-loop one? In fact, isn't that pretty much what we do, although informally? We're trying to fit a linear model to the data to obtain the minimum least-squares error of prediction, aren't we? The only difference is that our linear model embodies a closed loop.

You've had a lot of experience with statistics; you even wrote a book on it. Do you think you could take the same basic mathematical methods and alter them for use with a control-system model?

Greg Williams (921012-2) --

>And I suppose that there are more than a single scale of
>upsetnesses -- some upsetnesses occurring with reorganization, some
>not. One kind of upsetness (which can be more or less in amount)
>can occur when the success of controlling is in doubt, but
>reorganization isn't triggered.

When success of controlling (say, for exiting a theater when someone yells "Fire!") is in doubt, what doubts it? I think you need a hierarchical model -- if you could bring yourself to consider it as more than a loose and unimportant aspect of PCT.

>If you don't think this sort of upsetness is reasonable to
>postulate, please consider again the example of successfully
>exiting a theater after "Fire!" has been yelled, without (we both
>apparently agree) reorganization; it is unlikely, I believe, that

>the exiter would not have been at least a bit upset during the
>exiting.

As I have modeled reorganization, the outcome of reorganization is a control system that acts to prevent critical error by efficient nonrandom control of something that would otherwise cause critical error. In this example, the controlled variable might be something like a perception of oneself inside a building that's on fire. Getting out of the building -- reducing this perception to a reference level of zero -- might take some time; you wouldn't want reorganization to kick in when the control system is working as well as possible. So the reorganizing system has to work more slowly than the learned control system works. The "upsetness" you feel while exiting but not yet outside may reflect the beginnings of an internal disturbance, but before this disturbance can cause reorganization to start, you have acted and the fear, etc., has subsided.

To see how the very same small upset might become a very large one, just imagine that you step over to the exit door and find it locked or welded shut, while still believing that you're inside a building that is on fire. The difference is quantitative, not qualitative.

>I do think, though, that it might be possible for acquired criteria
>to override the built-in ones. I don't think that possibility is a
>problem for either of our viewpoints.

I agree in general, but I'm not sure what operation you mean by the term "override." The reorganizing system doesn't want any particular behavior to happen; its action is to alter organization, not to create a particular behavior. An acquired system is organized to control a particular variable. There can't be any conflict between reorganizing and systematically controlling. If an acquired system uses an action or pursues a goal that increases critical error (which we both agree is quite possible), this will simply cause reorganization to start. The reorganizing system has no direct way of opposing the control actions of an acquired system; it does not even know what they are. The critical error might be corrected if some other system is reorganized to conflict with the system producing the error, crippling it (actually, both). This could result in an overall reduction in critical error. The reorganizing system is not intelligent or foresighted. It simply keeps working toward a state of least critical error -- zero, ideally. That state is not reached in many people, or for long.

>>But how about the reorganization needed to perceive me as a crook
>>instead of how you were perceiving me before you got the new information?

>Or how about reorganizing to perceive that theater as a fire trap?
>Again, I'll be happy with consistency either way: exiting and con-
>realizing WITHOUT ANY reorganization, or both WITH SOME (perhaps
>minimal) reorganization. My problem with a lack of consistency is
>that you seem to want to treat the two examples as fundmentally
>different, yet I don't see a fundamental difference.

Maybe the discussion above removes some of the apparent inconsistency. The control hierarchy is learned specifically as a means of preventing critical error from becoming large enough to cause significant reorganization. That's automatic; reorganization simply continues until the critical error IS prevented from becoming that large. When the learned control processes work well enough, critical error does not become large enough to cause their organization to be altered. That's why they persist.

I do not, by the way, equate sensory experience of bodily states with critical error. Such sensory experiences -- of emotions, for example -- belong in the learned hierarchy. But they become, through reorganization, indicators of inner states that are learned to be "bad" or "good." The reorganizing system must work before such sensed inner states acquire any meaning. When you feel fear in the building on fire, this reflects a state of bodily preparedness for action, together with an error in the system that's trying to get you out of there. The reorganizing system, I would assume, treats a protracted state of bodily preparedness for action (without action to use up the energy) as a critical error. But the reorganizing system does not feel fear. It must know that this state is to be avoided before the learned system becomes able to sense it and treat it as a perception of "fear" to be reduced to zero.

>If memory, not reorganization, is involved in a particular instance
>of "facilitation" (or, more generally, "purposive influence"), then

>that instance is a kind of "rubber-banding," which might not have,
>for you, what you call "deep theoretical significance," but
>certainly has great practical significance AND scientific
>significance, in my opinion.

How is it an example of "rubber-banding?" I don't understand.

>>I wish you wouldn't use "controlling for" in this loose way, when
>>what you mean is "wishes to see."

>This is the crux of our dispute. I claim that this is truly
>CONTROLLING FOR, not just "wishes to see." B arranges A's
>environment so as to encourage a class of actions by A which B
>wants to see. If A doesn't perform actions in the class defined by
>B, then B RE-arranges A's environment. And so on, until A does
>actions in the class defined by B, or B gives up.

If this is the crux of our dispute, then our dispute seems to come down to a quantitative question: loop gain. I guess I automatically dismiss examples in which the loop gain is so low that disturbances can't be significantly resisted. A model of the sort of situation you propose just above would, I imagine, have a loop gain very much less than -1; the degree of control possible would be very low. For significant control, I use a rough rule of thumb of a loop gain of at least -5 or -10. Only when the loop gain becomes that large do you begin to see the typical properties of a control system -- action opposing disturbance, controlled variable remaining near the reference level.

Don't get me wrong; I'm not saying that people can't TRY to control others by means like the one you suggest. I'm not even saying that they don't convince themselves that they ARE controlling others by such means. But whatever control does exist is mostly in the imagination. Just consider the looseness and uncertainty in the scenario you propose. The would-be controller "encourages" a "class" of behaviors. The other person may or may not produce something in that class. If not, the controller tries a rearrangement of the environment and looks again to see if the desired outcome has happened, and so on until either it happens or the controller gives up and admits a lack of effect. If any sort of disturbance occurs that calls for the controllee to focus on behaviors of a different class, how much effect can the controller have in restoring the behaviors to the class the controller wants to see? The controller's effects are small, statistical, unreliable, and exceedingly slow. The loop gain must be close to zero. Think how easy it would be for the putative controllee to see the point of what the controller is doing and simply decide not to cooperate. Always assuming, of course, that there is no underlying threat of irresistable force that itself would be the actual means of control.

>IN PRACTICE, I see that this works much of the time: A indeed does
>perform actions B wants to see, and often within a short time.

If that is what you see, the only explanation I can think of is that you have misconstrued what you see. A much simpler explanation is that A has perceived what all of B's elaborate preparations are aimed at, and has decided to help B out by doing what B wants. I could see that as leading quickly and specifically to production of the behavior that B wants to see. Of course B might take this to indicate success of his or her method of control, particularly if it's control of A that B wants. I don't see how the method that you have outlined could be either quick or specific. Perhaps you have left something out of the description.

>In principle, there is no difference between this sort of control
>and the control of a cursor subject to a "hidden" disturbance -- in
>both cases, what is tending to thwart control cannot be "seen."

Qualitatively, perhaps not. But control is not just a qualitative matter -- control or no control. A control system that can cancel only 10 percent of a disturbance isn't much of a control system. In a tracking situation, 98 percent of the disturbance is cancelled.

>However, our models of physics suggest that wild fluctuations in
>the gravitational constant, beginning at 2 PM today, are unlikely.
>And -- here is where Skinner feared to tread -- PCT models suggest
>what constraints are important in determining the likely success or
>failure of attempts at "purposive influencing."

Let's leave physics out of this. The inanimate world is highly predictable and doesn't require much effort to control. You don't need much effort when all you have to do is set up initial conditions and let the physical system play out the consequences to the predicted (and wanted) end.

As to constraints, I agree. But let's not forget the constraint that you need some minimal amount of loop gain in order to see any important degree of purposive behavior.

>I suppose "innocuous" is in the eye of the beholder. Exchange relations
>seem rather innocuous to me, but maybe I'm just not enough of a revolutionary.
>Most of the time, I don't mind not being able to spend other people's money.
>But some people do mind that "imposition," much of the time -- I realize that.
>I'm not a Pollyana: NOT ALL social interactions are "win-win". But I don't
>think all are "lose-lose" or "win-lose," either.

Exchange relations are not control of another person. They specifically avoid the abitrary influencing of one person's actions to satisfy the goals of another. One person does not study another simply to get what is wanted out of the other; that is a control relation. Instead, each person considers what he or she has to offer that the other might want, and that is not inconvenient to give. If this is the understood basis for social interactions, then one doesn't need to manipulate others, because they will be doing the same thing. A simple request will suffice to obtain what you need that you can't get for yourself -- if not from one person, then from another. Often, simply the fact that you're having difficulty with a control problem will be enough to attract aid. And of course, a simple request from someone else will suffice for you to offer what is wanted, if not inconvenient to you. That's the system to which most people would subscribe under that kind of understanding of the social system.

This is a very different social relationship from one in which people memorize each others' characteristics, plot and intrigue, manipulate situations and environments, all so they can get what they want even if the other person doesn't want to cooperate or doesn't know there is manipulation going on. This latter kind of social organization is the one we have now -- when it's working at its best. Even at its best, it doesn't work very well. There is constant risk of conflict and escalation to violence. It's difficult to get what you want or need from other people, because everyone is defensive about "being controlled." They're defensive about that because that's what THEY are trying to do; they want to be the controller, not the controlled. Controlling for what you want is difficult because there's no simple way to get it when others are involved. The loop gain isn't very high. Often it's vanishingly low, but the desire to be in control makes people delude themselves that their efforts are actually working -- one wouldn't go to all that trouble for nothing, would one?

>>The con man can be sure of fooling the mark if he can try
>>his pitch on as many people as he likes and count only the
>>successes.

>The big-con artists do not operate on a statistical basis. They take time
>and pains to model the control structure of each potential mark, and give up
>(as PCT suggests they should) if the mark doesn't appear to want what they
>need the mark to want, in order for their (the con artists') controlling,
>which depends on the mark's actions, to be successful.

Why isn't that a statistical basis? You try a lot of possible cases, and sieve out the probables. This improves your chances, to be sure. The big-time con man looks for people who are asking to be conned, and bets that he's reading them right. All things considered, I wonder what the hourly pay of the average big-time con-man is. It's probably better than minimum wage, but not much. It's probably about the same as for anyone who lives by trying to hit it big. I've heard that the average thief lives in poverty. It's just that "living free" is more interesting than going straight. If you don't count jail time, which they don't.

I don't think that big-time con men constitute a significant fraction of the population. They don't cause the social and psychological problems of the world. They just take advantage of them, like carrion eaters. Even a lion doesn't need to know control theory to pick out the weakest members of the herd. Neither does a vulture.

From my point of view, the best use of control theory would be to strengthen the herd.

>Most of the methods you propose for controlling other people, or
>even predicting their behavior, simply won't work in the wild.

I disagree. I see them working "in the wild." (Yes, even AWAY from wild Black Lick Hollow.)

And I claim that you're misinterpreting what you see -- especially the part where you see them "working." Try a different interpretation.

>Today my son Evan was having a problem with his new birthday present, a
>radio- controlled truck. He asked me to help him figure out what was wrong
>with the transmitter. Some experiments guided by me showed a weak battery.
>Next time he'll be able to cure the malady himself. No, he didn't hold a
>gun to MY head, either. We BOTH got to where we wanted to be.

See how easy it is when nobody is trying to figure out how to control someone else? He asks, you give. The hardest part for you is waiting to give until he asks.

Curt McNamara (921013) --

>However, it is our actions that are judged and endure in the world, not our
>control structure nor purposes. In addition, actions are the visible
>(tip of the iceberg) portion of our control structure. So the crux of my
>argument is that it is (perhaps unintentional) actions (byproducts of
>control) which drive the world.

I think we also judge people by their intentions. You know, "Why are you being so nice to me today?"

And of course if it were not for intentions successfully achieved, the world would be a pretty random place.

Best to all, Bill P.

Date: Tue Oct 13, 1992 4:13 pm PST
Subject: nada

to dennis delprato: i have recieved your paper
to w.t. bourbon: since michelle's proposal is next friday, please bring the "to control or to be controlled" paper

i.n.kurtzer

Date: Wed Oct 14, 1992 3:12 am PST
From Greg Williams (921014)

>Bill Powers (921013.0930)

GW>>I do think, though, that it might be possible for acquired criteria
GW>>to override the built-in ones. I don't think that possibility is a
GW>>problem for either of our viewpoints.

>I agree in general, but I'm not sure what operation you mean by the
>term "override."

"Override" means that the error relative to an acquired reference signal, necessary for reorganization to start, is less than the errors relative to inherited reference signals, necessary for reorganization to start in the absence of the acquired reference signal.

>The control hierarchy is learned specifically as a means of preventing
>critical error from becoming large enough to cause significant
>reorganization. That's automatic; reorganization simply continues
>until the critical error IS prevented from becoming that large. When
>the learned control processes work well enough, critical error does
>not become large enough to cause their organization to be altered.
>That's why they persist.

Sounds good to me.

GW>>If memory, not reorganization, is involved in a particular instance
GW>>of "facilitation" (or, more generally, "purposive influence"), then
GW>>that instance is a kind of "rubber-banding," which might not have,
GW>>for you, what you call "deep theoretical significance," but
GW>>certainly has great practical significance AND scientific
GW>>significance, in my opinion.

>How is it an example of "rubber-banding?" I don't understand.

If the control structure stays the same (no reorganization), then new information simply disturbs a controlled variable, resulting in "offsetting" actions to maintain control of that variable. If you are controlling for phoning someone and I tell you that the phone number was recently changed, you control by using you fingers to dial the NEW number -- different actions, same controlled variable.

>If this is the crux of our dispute, then our dispute seems to come
>down to a quantitative question: loop gain.

I think the loop gain can range from very low to very high for a swimming teacher (as one instance), just as the loop gain can range widely for a subject controlling (in Rick's famous experiment) for keeping a dot near a certain point on a computer screen when the dot is subject to a random alteration in the direction of its movement each time the subject presses a key. I think the situations are analogous.

>Don't get me wrong; I'm not saying that people can't TRY to control
>others by means like the one you suggest. I'm not even saying that
>they don't convince themselves that they ARE controlling others by
>such means. But whatever control does exist is mostly in the
>imagination. Just consider the looseness and uncertainty in the
>scenario you propose.

"Looseness" and "uncertainty" need to be evaluated by looking at whether this kind of control works virtually all the time or only part of the time or never. As I look at the "wild," it works quite efficiently virtually all the time if the controller has an accurate model of something the other wants.

>The controller's effects are small, statistical, unreliable, and exceedingly
>slow. The loop gain must be close to zero.

Just as in Rick's experiment, the effects needn't be slow, unreliable, exceedingly slow, OR statistical (well, you might use statistics to economically describe the TRAJECTORIES in both cases, but the controller's intended outcome, namely, seeing the cursor or other person's actions he/she wants to see, does NOT need a statistical description, only a criterion (set by the controller) for meeting/not meeting the desire, which can be met by a CLASS of POSSIBLE actions). Aside: now that I think about it, there is no need for the class of possible actions to be finite; it is a CONCEPTUAL class, i.e., "the infinite class of all possible trajectories of the dot which remain within an inch of the dot," or "the infinite class of all possible ways of coming to stay afloat without external support in deep water."

>Think how easy it would be for the putative controllee to see the point of
>what the controller is doing and simply decide not to cooperate. Always
>assuming, of course, that there is no underlying threat of irresistable force
>that itself would be the actual means of control.

SEZ PCT: THE CONSTRAINT ON SUCCESSFUL CONTROL OF YOUR PERCEPTIONS WHICH DEPEND ON ANOTHER'S ACTIONS IS THAT YOU DISTURB (OR CAREFULLY DON'T DISTURB) IN SUCH A WAY THAT THE OTHER'S ONGOING CONTROL (IF THERE IS NO REORGANIZATION) OR NEW CONTROL (IF THERE IS REORGANIZATION) RESULTS IN WHAT _BOTH_ PARTIES WANT.

To meet this constraint requires interacting with the other enough (maybe just statistically, at the population level) to make a model of part of the other's control structure. If the model is poor, the other might "simply decide not to cooperate" -- but if the model is good, and the putative controllee sees the point FOR THE CONTROLLEE of what the controller is doing, then the controllee WILL cooperate, because cooperation will get the controllee what he/she wants. (Of course, if deception is involved, or if the controllee is unsophisticated, the controllee might only BELIEVE that he/she will get what he/she wants and not get what he/she doesn't want. Folks who understand the very real possibility of this sort of thing "patronize" others by warning them that what they (currently) want might not turn out to be best for them.)

GW>>IN PRACTICE, I see that this works much of the time: A indeed does
GW>>perform actions B wants to see, and often within a short time.

>If that is what you see, the only explanation I can think of is that
>you have misconstrued what you see. A much simpler explanation is that
>A has perceived what all of B's elaborate preparations are aimed at,
>and has decided to help B out by doing what B wants. I could see that
>as leading quickly and specifically to production of the behavior that
>B wants to see.

You've almost got it. Actually, B wants A to "help out" A, which happens (perhaps unbeknownst to A) to "help out" B, and might (but need not and usually doesn't!!!) in fact (via 20-20 hindsight) result in what A wouldn't consider beneficial for him/her.

>Of course B might take this to indicate success of his or her method of
>control, particularly if it's control of A that B wants.

Yes, and it would be "success" in the sense that B could have made a POOR model of what A wants, and thus would have failed. But B succeeded, because B met the constraints for successful control of his/her perceptions which depend on some of A's actions.

>I don't see how the method that you have outlined could be either quick or
>specific. Perhaps you have left something out of the description.

I hope the above helps. Perhaps the analogy with Rick's experiment will be most instructive for you.

>Exchange relations are not control of another person. They
>specifically avoid the abitrary influencing of one person's actions to
>satisfy the goals of another. One person does not study another simply
>to get what is wanted out of the other; that is a control relation.

"Arbitrary" is in the eye of the beholder. I claim that attempted control of ANYTHING is subject to certain constraints (you can't lift a two-ton rock by hand, and you can't make a person want what they don't want). I've been saying all along that I want to look at the CONSTRAINTS ON CONTROL of one's perceptions which depend on others' actions. IN GENERAL, CONTROL CANNOT BE ARBITRARY. If what you mean by "control of others" is ARBITRARILY MAKING THEM ACT AS YOU WISH, then you're NOT talking about what I am talking about (except in number 4 of my summary -- using force/threat of force) -- and even that kind of control cannot be ABSOLUTELY arbitrary (sticking a gun to my head won't result in my picking up a two-ton rock for you, sorry, better shoot).

>This is a very different social relationship from one in which people
>memorize each others' characteristics, plot and intrigue, manipulate
>situations and environments, all so they can get what they want even
>if the other person doesn't want to cooperate or doesn't know there is
>manipulation going on.

(Common) exchange interactions as well as the (not so common) nasty stuff you cite here BOTH involve attempts to control one's perceptions which depend on others' actions, and so should fit somewhere in my summary schema (the four kinds of "social" control). And PCT explains why exchange is more common than the nasty stuff.

>It's difficult to get what you want or need from other people, because
>everyone is defensive about "being controlled."

Some definitely with much higher loop gain than others. But the other's loop gain regarding "not being controlled" isn't a problem for a controller with a good model of what the other wants. In types 1-3 of my summary, the controllee is getting what he/she wants (at the time) for control to be successful.

>They're defensive about that because that's what THEY are trying to do; they

>want to be the controller, not the controlled.

Some definitely with much higher loop gain than others. The success of types 1-3 control requires that the controllee feel "in control," not "being controlled."

GW>>Today my son Evan was having a problem with his new birthday
GW>>present, a radio- controlled truck. He asked me to help him figure
GW>>out what was wrong with the transmitter. Some experiments guided by
GW>>me showed a weak battery. Next time he'll be able to cure the
GW>>malady himself. No, he didn't hold a gun to MY head, either. We
GW>>BOTH got to where we wanted to be.

>See how easy it is when nobody is trying to figure out how to control
>someone else? He asks, you give. The hardest part for you is waiting
>to give until he asks.

You seem to equate ALL of control of one's perception which depend on some of another's actions with (mainly) type 4 in my summary and some sub- types (involving deception and unsophisicated controllees) of type 2. I believe those types/sub-types are rare relative to the other types/sub-types in everyday life. Much control is (approximately) symmetric: win-win. I have yet to be convinced that ANY social interactions involving intention on the part of at least one of the parties involved do NOT involve one of the four types of control of one's perceptions which depend on others' actions. Such interactions include "exchange" as well as "teaching" and "con games."

I believe that an understanding of the constraints we face in attempting control of types 1 through 3 might help to avoid escalation to type 4.

>I think we also judge people by their intentions. You know, "Why are
>you being so nice to me today?"

We certainly do. (And not always cynically. "His heart is in the right place.") But believing it important to CHANGE their intentions is a sure path to violence. Attempting to see them ACT the way you want need not be -- if you take their intentions into consideration. Sez PCT.

Best (now why did he say that?),

Greg

Date: Wed Oct 14, 1992 7:04 pm PST
Subject: echo, echo echo...

to b.powers: a student here at SFASU would like to subscribe to the net and i forgot the subscribe command (the specifics) so could you please send that to me, thankyou.

also i sent a message to martin about his position on the predictions of pct, since he hasn't responded and since i felt what i said was worth at least saying it was wrong/misguided/correct/whatever please give my some imput (that includes anyone else if your willing to respond to the babbles of a grunt-student feel free). the message was sent about two/three? days ago but here's the basic point:

PCT can predict at a distance (i.e. into the future significantly further than the immediate). it is highly predictable whether you will sucessfully drive to the store without wrecking the car, or killing yourself or other people --given that is your intention. however, immediate predictions (actions carried out to oppose disturbances) will be known only if you know "other things" : certain system characteristics (ex. gain), and the vector sum of all disturbances at a given time (which could be inferred by the actions but that is circular), and the so-called "initial conditions" (this point can be and probably will\ be thrown away). that is a lot of things to know, especially if you are not the manipulator of them all which probably isn't the case. i realize that some of the disturbances may be remote but 1) there effects aren't 2) by definition if it is a disturbance it disturbs. this, in my opinion, is a rewording of stable end by various means.

thankyou from a grunt i.n.kurtzer

Date: Thu Oct 15, 1992 3:10 am PST
Subject: \... Echo!

From Greg Williams (921015)   >Isaac Kurtzer (921014)

Hi Isaac, from the Family Man with the longest hair in the room. It would be helpful to others on the net if you put your name and date at the beginning of each of your posts (something similar to the format I used at the beginning of this post would be nice; most netters have adopted Gary Cziko's suggestion of date code as YRMODY in numbers). It makes it easier to keep track of who replied to what. Thank you.

>PCT can predict at a distance (i.e. into the future significantly further
>than the immediate). it is highly predictable whether you will sucessfully
>drive to the store without wrecking the car, or killing yourself or
>other people --given that is your intention. however, immediate
>predictions (actions carried out to oppose disturbances) will be known
>only if you know "other things" : certain system characteristics
>(ex. gain), and the vector sum of all disturbances at a given time
>(which could be inferred by the actions but that is circular), and
>the so-called "initial conditions" (this point can be and probably will\
>be thrown away). that is a lot of things to know, especially if you
>are not the manipulator of them all which probably isn't the case. i
>realize that some of the disturbances may be remote but 1) there effects
>aren't 2) by definition if it is a disturbance it disturbs. this, in my
>opinion, is a rewording of stable end by various means.

Knowing an organism's controlled variable allows the knower to predict the desired outcome which is sought by the organism. I.e., knowing that a rat is hungry, because you have kept it away from food for several hours, you can predict the OUTCOME when the rat is provided with food: it will eat.

But predicting the PARTICULAR actions (outputs) used by the hungry rat to satisfy its desire for food if it is provided -- that is, the PARTICULAR actions the rat uses to eat -- is, as you say, difficult. What is often possible in practice (ranging all the way from cases involving Skinner boxes to everyday "wild" human behavior) is, given good guesses about an organism's desired outcome and knowledge about some of the organism's other characteristics and about some of the characteristics of the organism's environment, to be able to predict that the organism will use actions in a certain class of possible actions to attempt to achieve the desired outcome.

A hungry rat in a Skinner box with a lever which releases food will perform an action in the class of "moving the lever." The Porsche driver approaching a left-hand turn which he wants to negotiate successfully will perform an action in the class of "turning the steering wheel counterclockwise." The size of the predicted class of actions will vary from situation to situation.

I claim that often in social interactions, party A, who is controlling for some perceptions which depend on actions of party B, can successfully control IF B's actions, upon which A's controlled perceptions depend, are in a certain class of actions (the class "desired" by A).

Furthermore, I claim that often in social interactions, party A can successfully arrange (disturb or carefully not disturb) B's environment so that party B indeed performs actions in that certain class of actions.

So, to control his/her perceptions depending on B's actions, A needn't predict B's actions EXACTLY, and (on the basis of observations) I claim that such control is often achieved.

From a fellow student grunt (I'm always learning!),   Greg

Date: Thu Oct 15, 1992 7:26 am PST
Subject: Put your model where your mouth is.

[From Bill Powers (921015.0700)]   Greg Williams (921014) --

>"Override" means that the error relative to an acquired reference signal,
>necessary for reorganization to start, is less than the errors relative to
>inherited reference signals, necessary for reorganization to start in the
>absence of the acquired reference signal.

Still can't figure out what you mean. It would help if you described the specific situation you have in mind instead of the generalization you got from it. A diagram would help even more.

RE: new information as rubber-banding.

>If you are controlling for phoning someone and I tell you that the
>phone number was recently changed, you control by using you fingers
>to dial the NEW number -- different actions, same controlled
>variable.

The actions are different at the level of moving your fingers, but the same at the level of phoning someone. The plan "Call Joe" remains the same, and the perception matches it, at the higher level. What has to change is the way the error at that level is translated into a specific sequence of digits to dial to correct the error. By asking for Joe's phone number, I obtain a new number in my memory. This does not make it a reference signal yet. This new number must be selected when I want to achieve "Call Joe." This means that the higher system must, given the same error signal as before, select a new reference- sequence at the lower level. This requires reorganization. If you have been calling Joe from a memorized phone number for years, chances are that the first few times you try to call after the number has changed, you will "forget" that it's been changed and call the old number. It takes a few errors to make the new number "stick" -- i.e., to get the connection to the new reference signal changed.

This involves reorganization, but perhaps not carried out in the way I visualize for "E. coli" learning. This sort of phenomenon might be a clue about a more systematic way of reorganizing. But it still takes some time to make the old connection go away and for the new one to become as automatic as the old one was. You may have to stop and consciously "remind yourself" that Joe has a new number. And as someone remarked on the net once, this is likely to result in a repetitive sequence of starting to dial the number, reminding yourself that it's changed, and dialing the new number; it may take a long time to get to dialing the new number without going through that sequence. This is probably a clue, too.

To avoid going through this sort of reorganization, people write down phone numbers in a book; erase the old number and put the new one in. Then nothing has to reorganize. They just look up Joe's number and call it.

>>If this is the crux of our dispute, then our dispute seems to come
>>down to a quantitative question: loop gain.

>I think the loop gain can range from very low to very high for a
>swimming teacher (as one instance), just as the loop gain can range
>widely for a subject controlling (in Rick's famous experiment) for
>keeping a dot near a certain point on a computer screen when the >dot
is subject to a random alteration in the direction of its >movement
each time the subject presses a key. I think the >situations are
analogous.

I'm not sure what the loop gain of E. coli would be if E. coli's random actions had to control another E. coli's swimming behavior by disturbing the other's time rate of change of concentration in order to control the first E. coli's sensed time rate of concentration. The thought of a teacher randomly trying different teaching methods as a way of helping a student randomly reorganize toward a specific behavior does not impress me as fraught with possibilities.

Your proposal calls at least for some experimental or working-model support.

>"Looseness" and "uncertainty" need to be evaluated by looking at
>whether this kind of control works virtually all the time or only
>part of the time or never. As I look at the "wild," it works quite
>efficiently virtually all the time if the controller has an >accurate
model of something the other wants.

I think it's time for evidence in the form of examples, and some backing up of the generalities by showing a model that would work as you suggest.

>>The controller's effects are small, statistical, unreliable, and
>>exceedingly slow. The loop gain must be close to zero.

>Just as in Rick's experiment, the effects needn't be slow,
>unreliable, exceedingly slow, OR statistical ...

But Rick's experiment had to do with an organism reorganizing to control one of its OWN critical variables, not one organism trying to use another one doing the same thing to achieve the first organism's goals. What we need is an experiment with Rick's model in which a human teacher tries to "facilitate" E. coli's progress up the gradient. Put your model where your mouth is.

>Aside: now that I think about it, there is no need for the class of possible
>actions to be finite; it is a CONCEPTUAL class, i.e., "the infinite class of
>all possible trajectories of the dot which remain within an inch of the dot,
>" or "the infinite class of all possible ways of coming to stay afloat
>without external support in deep water."

This is illustrates a problem with arguing at high levels of abstraction: general statements end up saying much less than they seem to say. "The infinite class of all possible ways of coming to stay afloat without external support in deep water" can be stated much more succinctly: you're trying to describe "swimming." That is the consequence of the actions that you're trying to get the students to perform. You are simply describing the teacher's reference level for what is to be learned: the outcome that is to result. In effect, you're saying "I can't teach them the actions that will result in swimming, but by George I'll know swimming when I see it." To predict that they will then be doing one or more of the things that result in staying afloat without external support will not be of much use while they're trying to find some action in that class, and perform it so it has the desired effect.

You're trying to weasel out of the difficulties in exactly the way Skinner did. A response is that class of actions that has a particular consequence. All Skinner did was formalize (vaguely) the same habit that all behavioral scientists follow: naming behaviors by their controlled consequences, so as not to have to explain how the organism could select behaviors that, combined with disturbances, end up producing the same consequences again and again.

>SEZ PCT: THE CONSTRAINT ON SUCCESSFUL CONTROL OF YOUR PERCEPTIONS
>WHICH DEPEND ON ANOTHER'S ACTIONS IS THAT YOU DISTURB (OR CAREFULLY
>DON'T DISTURB) IN SUCH A WAY THAT THE OTHER'S ONGOING CONTROL (IF
>THERE IS NO REORGANIZATION) OR NEW CONTROL (IF THERE IS
>REORGANIZATION) RESULTS IN WHAT _BOTH_ PARTIES WANT.

This, too, is merely explaining the outcome by describing the outcome. If both parties do end up controlling successfully, then the actions each takes to counteract the disturbances from the other (or from any source) cause no important errors in either of them. If that's not true, there will be conflict. It's not necessary for either control system to refer to this abstract constraint as they learn to interact with each other. Each system will either meet with no resistance, or have its actions resisted. If its actions are seriously resisted, and there's no alternative already known, the organism will begin to reorganize. If its actions aren't resisted, the organism will simply control. It doesn't have to know that it's using the actions of the other as part of its control loop. That doesn't require any planning; it simply happens, if it happens. If it doesn't happen, that's OK, too: control occurs either way. The main thing is to control your perceptions; the actions by which you accomplish that will come to be whatever is required. In making your way to the other side of a crowded room, your actions will probably involve the actions of many other people; in an empty room the same goal will be achieved without anyone else. You don't have to predict how each person in the crowded room will react to your disturbances. You just make your way, muttering "Pardon me, sorry, oops, pardon me" -- or you just put your head down and push. If someone won't get out of your way you find a different way through.

There's too much abstract conjecture going on here. Let's try to tie this argument to specific examples. Better yet, let's stop fooling around in the stratosphere, and start proposing some experiments to test all these deductions and pseudo-deductions. We sound like a couple of psychologists or philosophers. Let's get back to science.

Best,   Bill P.

Date: Thu Oct 15, 1992 2:24 pm PST
Subject: Apparent S-R behaviour: what goes on?

[Martin Taylor 921015 16:00]

The following question probably indicates that my understanding of PCT is less secure than I have been believing it to be, but here goes, anyway.

On two occasions recently. I observed what seems on the face of it to be S-R behaviour--the situational context causing a non-functional behaviour to be executed. On one of those occasions, I was the actor, and it felt as if this behaviour was "extracted" automatically from me rather than being the control of any perception of which I was then (or now) aware. Here's the situation.

Event 1: I was carrying a door upstairs in my house. The stairs have rather nice panelling on the side. As I turned the corner at the top, the swing of the door caused it to bump the panelling lightly. I immediately said "Sorry" and then wondered why I did so. As a conditioned response in the S-R tradition, it's easy to understand. But where does it fit in PCT?

Event 2 is very similar. I was approaching a double door that has glass windows as its upper panels. When I was about 5 yards away, a woman came through in the other direction. When she opened the door, through which she could easily see anyone on my side, the door bumped a cardboard box she had not sees, left by a cleaner or somebody. She also said "Sorry" immediately, even though she was well aware she had not hit anyone.

The trigger for writing this note was Bill's comment on dialing the old phone number for someone whose number recently changed. It's not quite the same situation, but there seems to be something in common.

Bill, or anyone: is there a straightforward PCT interpretation of this immediate execution of "Sorry" following a bump when one is well aware that there is no-one there to accept the apology. I imagine that this situation has many analogues in other contexts, and it is a bit puzzling at present.

Martin

Date: Thu Oct 15, 1992 3:10 pm PST
Subject: family man

to g. cziko: didn't mean to override you if that was your impression, b.powers helped me out last time that's all.   thankyou

to w.t.bourbon: yes it is ISAAC; am anxious to recieve your papers; see you oct. 30 thankyou

to g.williams: i appriciate the response.   thankyou

to r."the player" marken: in case you forgot, i WOULD like a copy of the recently rejected paper. y thankyou

to d.delprato: reading your article am i to understand that skinner felt that psychology could be reduced to chemistry, physics, etc. but that it was not necessary to explain behavior (superfluous?) ?

Silly snorts from a garden-variety grunt   i.n.kurtzer

Date: Thu Oct 15, 1992 4:23 pm PST

Subject: Sorry phenomenon; social dogs

[From Bill Powers (921015.1800)] Martin Taylor (921015.1600) --

It would be interesting to compare an S-R explanation of the "Sorry" phenomenon with a CT explanation. Both would necessarily rely on some assumptions about other processes in the person, so we couldn't really decide which is "right." But it would be interesting to see what kinds of assumptions would be needed to make a coherent story. I think you have posed an excellent exercise for the student. How about some attempts from people who ordinarily don't say much on the net?

I'll set one ground rule: whatever explanation you offer on either side, it should include an experimental method for checking out each assumption (just a description, not an actual experiment).

In reply to direct post: I WOULD like to see the data from your experiment when you have time; I'll return it, so you don't need to go to the trouble of copying it.

Greg Williams (921015) --

Eavesdropping on your comments to Isaac Kurtzer:

>I claim that often in social interactions, party A, who is controlling for
>some perceptions which depend on actions of party B, can successfully control
>IF B's actions, upon which A's controlled perceptions depend, are in a certain
>class of actions (the class "desired" by A). Furthermore, I claim that often
>in social interactions, party A can successfully arrange (disturb or carefully
>not disturb) B's environment so that party B indeed performs actions in that
>certain class of actions. So, to control his/her perceptions depending on B's
>actions, A needn't predict B's actions EXACTLY, and (on the basis of
>observations) I claim that such control is often achieved.

A thought: I just realized that dogs can do the same thing with each other. But I presume that they don't know control theory and can't assess other dogs' goals. Can't a person accomplish control that brings another into the loop without having to know that the other person is also controlling something?

Best to all, Bill P.

Date: Fri Oct 16, 1992 12:21 am PST
From: Control Systems Group Network

[Oded Maler 921016] * [Martin Taylor 921015 16:00]

If you have an explanation of how saying "sorry" while bumping into a person works (btw, not in all cultures..) then the phenomenon you described can be explained by a misclassification of the complex perceptual variable "I just hit a person". Perceiving hitting is more elmentary the identifying personhood. When you are busy doing other tasks (your higher levels are occupied) you don't have time to to have a refined perecption of a non-person/person within the time-scale of the "sorry" loop, which is located rather low in your hierarchy.

Being infinitely fast means never having to say that you are sorry.

Sorry if it is not in exact but rather in qualitative PCT terms.

--Oded

Date: Fri Oct 16, 1992 6:15 am PST
Subject: Models and Data

From Greg Williams (921016) >Bill Powers (921015.0700)

GW>>"Override" means that the error relative to an acquired reference
GW>>signal, necessary for reorganization to start, is less than the
GW>>errors relative to inherited reference signals, necessary for

GW>>reorganization to start in the absence of the acquired reference
GW>>signal.

>Still can't figure out what you mean. It would help if you described
>the specific situation you have in mind instead of the generalization
>you got from it. A diagram would help even more.

It is simply YOUR reorganization model, except that the criteria for starting reorganization due to errors in (some? all?) innate critical variables are subject to alteration by reorganization triggered by sufficient error in acquired (within a single lifetime) reference levels -- which can be thought of as acquired critical variables. Consider simplified people with only one innate critical variable. The people walk around on a flat plain with edges; falling over an edge means death. Near the edge, the plain slopes downhill increasingly sharply. The innate critical variable is feet being angled (as they would be on a slope), so if a person is too close to an edge, it begins to reorganize. But now suppose the People of the Plain develop a people- sacrifice religion. The sacrificees-to-be are raised from birth to "believe" that, when the chosen (by the priests) time comes, their walking over the edge will win them the ultimate religious reward: eternal life. Their control structure reorganizes over their period of training so that an acquired critical variable having to do with what it takes to win eternal life "overrides" the innate critical variable having to do with angled feet (the criterion for critical error in foot angle for triggering reorganization gets larger), and, at the officious time, they walk right down the slope and over the edge, without any "random"-appearing movement (reorganization is not triggered by the angled-feet critical error, because the criterion for starting reorganization is not met -- well, until after the point of no return!). There they go, perhaps singing "holy" phrases.

>RE: new information as rubber-banding.

>The actions are different at the level of moving your fingers, but the
>same at the level of phoning someone. The plan "Call Joe" remains the
>same, and the perception matches it, at the higher level. What has to
>change is the way the error at that level is translated into a
>specific sequence of digits to dial to correct the error.

In rubber banding: The actions are different at the level of moving your [the subject's] fingers, but the same at the level of holding the knot over the dot. The perception "do as the experiment requested" remains the same, and the perception matches it, at the higher level. What has to change is the way the error at that level is translated into a specific sequence of positions of the [subject's] end of the rubber band.

>By asking for Joe's phone number, I obtain a new number in my memory.

Yes, unless you just write it down for future reference.

>This does not make it a reference signal yet. This new number must be selected
>when I want to achieve "Call Joe." This means that the higher system must,
>given the same error signal as before, select a new reference-sequence at the
>lower level. This requires reorganization.

Fine by me. But if you didn't have the old number memorized, and don't try to memorize the new one (you don't call Joe very often -- so seldom that you didn't even know he moved), this is just a kind of rubber-banding. That was all I was originally claiming. With memorization or re-memorization, I accept that there could be reorganization -- at least, some sort of learning; I wonder what critical variable might have reached its criterion error level, and it doesn't seem very "random" -- and so it would fall under type 3 in my summary of control of one's perceptions depending on another's actions.

>This involves reorganization, but perhaps not carried out in the way I
>visualize for "E. coli" learning. This sort of phenomenon might be a
>clue about a more systematic way of reorganizing.

I think you're headed in the right direction now, and not simply stochastically.

>I'm not sure what the loop gain of E. coli would be if E. coli's
>random actions had to control another E. coli's swimming behavior by
>disturbing the other's time rate of change of concentration in order
>to control the first E. coli's sensed time rate of concentration.

>The thought of a teacher randomly trying different teaching methods as a way
>of helping a student randomly reorganize toward a specific behavior does not
>impress me as fraught with possibilities.

I don't think E. coli is equipped to do this sort of thing. If you do, please present a model for how. The fallacy I see in your trying to analogize between E. coli and people is that the person who attempts to control his/her perceptions which depend on others' actions does NOT act randomly. That is NOT the way (successful) teachers teach. Ask Gary whether he teaches by "randomly trying different teaching methods"! Granted that Gary cannot predict in advance that student X will have a difficulty of kind y; still, Gary knows that when a student does (apparently, to Gary) have difficulty of kind y, then he will do z (because it has often worked before in such cases -- it might not work now, in which case he will try doing i; if i doesn't work, he'll just recommend that the student drop the course -- you can't teach some people statistics!).

>Your proposal calls at least for some experimental or working-model support.

I agree. But a working model of a teaching situation isn't possible at this time. PCT models are still quite primitive: arm movements and distance-seeking from obstacles. We're far, far, far from modeling the complex concepts involved in teaching. And I'm building a house. Nevertheless, I still can point to the abundant evidence that, in the real world, successful control of one's perceptions depending on others' actions is ubiquitous. People ARE "taught" to "swim." Can you FALSIFY my hypothesis that social interactions involving intention (that is, are non-"accidental") ALWAYS involve at least one of the four types of control in my summary? Or can you only suggest other possible hypotheses which (given current state-of-the-art) cannot be clearly regarded as more likely to be correct than my hypothesis? I think we're arguing in a vacuum -- a vacuum of data, not of models. If you can't falsify my hypothesis, you should look at "natural experiments" (the "wild") as well as design experiments to see which hypothesis best fits the data. (An aside: If ANYONE comes up with examples not fitting my types, I'd be happy to see the list expand.)

>What we need is an experiment with Rick's model in which a
>human teacher tries to "facilitate" E. coli's progress up the
>gradient. Put your model where your mouth is.

Before even attempting any models, you need to take the data on human social interactions seriously. A model of E. coli is NOT a model of a human. How about looking, instead, at how a human teacher "facilitates" a human student's "progress up the gradient" of learning some subject? I suspect that you'll object because of the complexities with respect to the current stage of PCT models, which are at the beginning stages currently. But it is better to look for the money near the dark spot where you dropped it than to look for it under a lamppost down the street. If a human teacher cannot "facilitate" E. coli's progress up a gradient, that tells me that the analogy between E. coli's klinotaxis and a human student's learning is a poor one. Human students CAN learn to swim -- and can often learn faster with a teacher's "facilitation." E. coli CANNOT learn to progress up a chemical gradient faster, regardless of "facilitation."

>This is illustrates a problem with arguing at high levels of
>abstraction: general statements end up saying much less than they seem
>to say.

So, let's get back to the data: data about human social interactions.

>In effect, you're saying "I can't teach them the actions that will result in
>swimming, but by George I'll know swimming when I see it."

I am saying that PCT says that the teacher can at least tell the student when his actions DON'T look like swimming actions, and suggest new, more appropriate (to learning how to swim, as judged by the teacher) lower-level reference signals for the student, rather than sitting around waiting for a purported random-walk process to achieve the student's desired goal ("swimming"). This is what Skinner called "shaping." (Skinner didn't care WHY it works!)

One way a human CAN "facilitate" E. coli's moving more quickly up a chemical gradient (I don't have to run the experiment, it obviously works, at least in a statistical sense of speeding progress up a gradient on average, over many trials): whenever E. coli starts to go in a direction with a component DOWN the gradient, almost immediately block progress, so the bacterium will just sit there, instead of going any farther down the gradient, before it tumbles. The problem with claiming any analogy between this and human teaching is that E. coli's gradient-climbing is never improved in the absence of the "barrier-facilitation," whereas the human student comes to be able to swim in the absence of the teacher.

>You're trying to weasel out of the difficulties in exactly the way
>Skinner did. A response is that class of actions that has a particular
>consequence. All Skinner did was formalize (vaguely) the same habit
>that all behavioral scientists follow: naming behaviors by their
>controlled consequences, so as not to have to explain how the organism
>could select behaviors that, combined with disturbances, end up
>producing the same consequences again and again.

I agree with you that how organisms produce invariable ends via variable means should be be explained. PCT models are the basis for my summary of the ways in which control of perceptions depending on others' actions can succeed. I am attempting to explain WHY Skinner's experiments succeeded in certain ways, and WHY they have limits to success. He didn't know why. But it is only fair to add that, for his purposes of "prediction and control" (AS HE SAW THEM -- I think he was fooling himself), he didn't think he needed to know why.

>This, too, is merely explaining the outcome by describing the outcome.
>If both parties do end up controlling successfully, then the actions
>each takes to counteract the disturbances from the other (or from any
>source) cause no important errors in either of them. If that's not
>true, there will be conflict. It's not necessary for either control
>system to refer to this abstract constraint as they learn to interact
>with each other. Each system will either meet with no resistance, or
>have its actions resisted. If its actions are seriously resisted, and
>there's no alternative already known, the organism will begin to
>reorganize. If its actions aren't resisted, the organism will simply
>control. It doesn't have to know that it's using the actions of the
>other as part of its control loop. That doesn't require any planning;
>it simply happens, if it happens. If it doesn't happen, that's OK,
>too: control occurs either way. The main thing is to control your
>perceptions; the actions by which you accomplish that will come to be
>whatever is required.

In all of this, some of the four types of control in my summary are happening. If the would-be controller has poor models of what the other wants, types 1, 2, and 3 won't succeed, and only type 4 is left -- conflict. If you don't care about understanding why some social interactions go (virtually) straight to 4, while others don't, that's OK. I know that others (especially those who would prefer to avoid type 4, if possible) DO care. Some individuals on the net have told me that they care.

>Let's get back to science.

Let's get back to data on human social interactions. When PCT models are up to dealing with that data, I'll be happy to aid with the modeling -- maybe I'll be done with house-building by then! Perhaps Tom Bourbon can provide some optimistic news on making PCT models for human social interactions. I hope so!

>Bill Powers (921015.1800)

>A thought: I just realized that dogs can do the same thing with each
>other. But I presume that they don't know control theory and can't
>assess other dogs' goals. Can't a person accomplish control that
>brings another into the loop without having to know that the other
>person is also controlling something?

That depends on what you mean by knowing another is controlling something. Most people don't know PCT (and no dogs do, presumably, and probably not even Bill's cat!), but (barring "coincidence"), successful control of perceptions which depend on another's actions (except for type 4 control) does indeed require having a model of (some of) what the other wants. Apparently, most PCT-ignorant people and at least some PCT-ignorant dogs (cats, too!) have such models, to a degree -- and it appears that dogs and cats have such models not only of other dogs/cats, but of humans and prey animals, too. Our dog

"Friendly" has come to make the model (I speculate) that Pat wants to go on a walk in the woods after watching a video movie (this occurs about once a week); Friendly is always there at the door with her tail wagging and her "please" look when the movie ends; she is not at the door, etc., otherwise. If this is anthropomorphizing, so be it. I don't follow Dr. Skinner's rules, because I think psychology can get farther along (even if we err to a degree toward over-anthropomorphizing in the short-term) by not following those rules and making tentative models of organisms' innards, which models must then be corrected in the light of experimental evidence. I hasten to add that it seems to me that E. coli has no "modeling others" capability -- I don't anthropomorphize THAT much! Many invertebrates -- much ignored by PCTers, to their great loss, I believe -- appear to have quite refined modeling capabilities, largely, I suppose, unmodifiable in a single lifetime.

I suppose Ed Ford could provide examples showing that explicit understanding of PCT (of which some people, but not other animals, are capable) can sometimes improve the ability to successfully control perceptions which depend on others' actions -- and even to do so without needing to resort to violence or threats of violence.

Best,   Greg

P.S. I'll mail the corrected figures on Monday. Sorry about my fumble on the function keys. Glad to hear the paper is finally ready to go, otherwise.

Date: Fri Oct 16, 1992 8:43 am PST
Subject: family man

Isaac

Here is the rejected paper. Enjoy.

The Hierarchical Behavior of Perception

Richard S. Marken
RSM & Associates
Los Angeles, CA 90024

Abstract

This paper argues that the coincidental development of hierarchical models of perception and behavior is not a coincidence. Perception and behavior are two sides of the same phenomenon -- control. A hierarchical control system model shows that evidence of hierarchical organization in behavior is also evidence of hierarchical organization in perception. Studies of the temporal limitations of behavior, for example, are shown to be consistent with studies of temporal limitations of perception. A surprising implication of the control model is that the perceptual limits are the basis of the behavioral limits. Action systems cannot produce controlled behavioral results faster than the rate at which these results can be perceived. Behavioral skill turns on the ability to control a hierarchy of perceptions, not actions.

Psychologists have developed hierarchical models of both perception (eg. Bryan and Harter, 1899; Palmer, 1977; Simon, 1972; Povel, 1981) and behavior (eg. Albus, 1981; Arbib, 1972; Greeno and Simon, 1974; Lashley, 1951; Martin, 1972; Keele, Cohen and Ivry, 1990; Rosenbaum, 1987). This could be a coincidence, a case of similar models being applied to two very different kinds of phenomena. On the other hand, it could reflect the existence of a common basis for both perception and behavior. This paper argues for the latter possibility, suggesting that perception and behavior are two sides of the same phenomenon -- control (Marken, 1988). Control is the means by which agents keep perceived aspects of their external environment in goal states (Powers, 1973). It is argued that the existence of hierarchical models of both perception and behavior is a result of looking at control from two different perspectives; that of the agent doing the controlling (the actor) and that of the agent watching control (the observer). Depending on the perspective, control can be seen as a perceptual or a behavioral phenomenon.

From the actor's perspective, control is a perceptual phenomenon. The actor is controlling his or her own perceptual experience, making it behave as desired. However, from the observer's perspective, control is a behavioral phenomenon. The actor appears to be controlling variable aspects of his or her behavior in relation to the environment. For example, from the perspective of a typist (the actor), typing involves the control of a dynamically changing set of kinesthetic, auditory and, perhaps, visual perceptions. If there were no perceptions there would be no typing. However, from the perspective of someone watching the typist (the observer), perception is irrelevant; the typist appears to be controlling the movements of his or her fingers in relation to the keys on a keyboard.

These two views of control have one thing in common; in both cases, control is seen in the behavior of perception. For the actor, control is seen in the behavior of his or her own perceptions. For the observer, control is seen in the behavior of his or her own perceptions of the actor's actions. (The observer can see the means of control but can only infer their perceptual consequences as experienced by the actor). If control is hierarchical then it can be described as the behavior of a hierarchy of perceptions. Hierarchical models of perception and behavior can then be seen as attempts to describe control from two different perspectives, those of the actor and observer, respectively. This paper presents evidence that hierarchical models of perception and behavior reflect the hierarchical structure of control.

A Perceptual Control Hierarchy

The concept of control as the behavior of perception can be understood in the context of a hierarchical control system model of behavioral organization (Powers, 1973; 1989). The model is shown in Figure 1. It consists of several levels of control systems (the figure shows four levels) with many control systems at each level (the figure shows seven). Each control system consists of an input transducer (I), comparator (C) and output transducer (O). The input transducer converts inputs from the environment or from systems lower in the hierarchy into a perceptual signal,p. The comparator computes the difference between the perceptual signal and a reference signal ,r. The output transducer amplifies and converts this difference into actions which affect the environment or become reference signals for lower level systems.

Insert Figure 1 Here

The control systems at each level of the hierarchy control perceptions of different aspects of the external environment. However, all systems control perceptions in the same way; by producing actions that reduce the discrepancy between actual and intended perceptions. Intended perceptions are specified by the reference signals to the control systems. The actions of the control systems coax perceptual signals into a match with reference signals via direct or indirect effects on the external environment. The actions of the lowest level control systems affect perceptions directly through the environment. The actions of higher level control systems affect perceptions indirectly by adjusting the reference inputs to lower level systems.

The hierarchy of control systems is a working model of purposeful behavior (Marken, 1986; 1990). The behavior of the hierarchy is purposeful inasmuch as each control system in the hierarchy works against any opposing forces in order to produce intended results. Opposing forces come from disturbances created by the environment as well as interfering effects caused by the actions of other control systems. The existence of disturbances means that a control system cannot reliably produce an intended result by selecting a particular action. Actions must vary to compensate for varying disturbances. Control systems solve this problem by specifying what results are to be perceived, not how these results are to be achieved. Control systems control perceptions, not actions. When set up correctly the control systems in the hierarchy vary their actions as necessary, compensating for unpredictable (and, often, undetectable) disturbances, in order to produce intended perceptions. Indeed, the term "control" refers to this process of producing intended perceptions in a disturbance prone environment.

Levels of Perception.

Powers (1990) has proposed that each level of the hierarchy of control systems controls a different class of perception. These classes represent progressively more abstract aspects of the external environment. The lowest level systems control perceptions that represent the intensity of environmental input. The next level controls sensations (such as a colors), which are functions of several different intensities. Going up from sensations there is control of configurations (combinations of sensations), transitions (temporal changes in configurations), events (sequences of changing configurations), relationships (logical, statistical, or causal co- variation between independent events), categories (class membership), sequences (unique orderings of lower order perceptions), programs (if-then contingencies between lower level perceptions), principles (a general rule that exists in the behavior of lower level perceptions) and system concepts (a particular set of principles exemplified by the states of many lower level perceptions; see Powers, 1989, pp. 190-208). These eleven classes of perception correspond to eleven levels of control systems in the hierarchical control model. All control systems at a particular level of the hierarchy control the same class of perception, though each system controls a slightly different exemplar of the class. Thus, all systems at a particular level may control configuration perceptions but each system controls a different configuration.

The rationale for hierarchical classes of perceptual control is based on the observation that certain types of perception depend on the existence of others. Higher level perceptions depend on (and, thus, are a function of) lower level perceptions. For example, the perception of a configuration, such as a face, depends on the existence of sensation (color) or intensity (black/white) perceptions. The face is a function of these sensations and intensities. The lower level perceptions are the independent variables in the function that computes the higher level perception. Their status as independent variables is confirmed by the fact that lower level perceptions can exist in the absence of the higher level perceptions, but not vice versa. Color and intensity perceptions can exist without the perception of a face (or any other configuration, for that matter) but there is no face without perceptions of intensity and/or color.

The Behavior of Perceptions. From the point of view of the hierarchical control model, "behaving" is a process of controlling perceptual experience. Any reasonably complex behavior involves the control of several levels of perception simultaneously . For example, when typing the word "hello", one controlled perception is the sequence of letters "h", "e", "l" ,"l" and "o". The perception of this sequence is controlled by producing a sequence of keypress event perceptions. Each keypress event is controlled by producing a particular set of transitions between finger configuration perceptions. Each finger configuration is controlled by a different set of force sensations which are themselves controlled by producing different combinations of intensities of tensions in a set of muscles.

The perceptions involved in typing "hello" are all being controlled simultaneously. Transitions between finger configurations are being controlled while the force sensations that produce the configuration perceptions are being controlled. However, the typist is usually not aware of the behavior of all these levels of perception. People ordinarily attend to the behavior of their perceptions at a high level of abstraction, ignoring the details. We attend to the fact that we are driving down the road and ignore the changing muscle tensions, arm configurations and steering wheel movements that produce this result. Paying attention to the details leads to a deterioration of performance; it is the opposite of "zen" behavior, where you just attend to the (perceptual) results that you intend to produce and let the required lower level perceptions take care of themselves (Herrigal, 1971). However, while it violates the principles of zen, attention to the detailed perceptions involved in the production of behavioral results can provide interesting hints about the nature of the perceptual control hierarchy.

The Perception of Behavior. The behavior of an actor who is organized like the hierarchical control model consists of changes in the values of variables in the actor's environment. An observer cannot see what is going on inside the actor; he or she can only see the actor's actions and the effect of these actions on the external environment. The effect of these actions is to cause purposeful behavior of certain variables in the environment; the variables that correspond to perceptions that the actor is actually controlling. The purposefulness of the behavior of these variables is evidenced by the fact that consistent behaviors are produced in the context of randomly changing environmental disturbances. Thus, a typist can consistently type the word "hello" despite changes in the position of the fingers relative to the keyboard, variations in the push-back force of the keys or even a shift from one keyboard arrangement to another (from QWERTY to Dvorak, for example).

Since the actor controls his or her own perceptions, the observer cannot actually see what the actor is "doing"; the actor's "doings" consist of changing the intended states of his or her own perceptions. All the observer sees is variable results of the actor's actions; results that may or may not be under control. For example, the observer, might notice that a click occurs each time the typist presses a key. The click is a result produced by the typist and the observer is likely to conclude that the typist is controlling the occurrence of the click. In fact, the click may be nothing more than a side effect of the typist's efforts to make the key feel like it has hit bottom. There are methods that make it possible for the observer to tell whether or not his or her perceptions of the actor's behavior correspond to the perceptions that are being controlled by the actor (Marken, 1989). These methods make it possible for the observer to determine what the actor is actually doing (i.e. controlling).

Hierarchical Control

The hierarchical nature of the processes that generate behavior would not be obvious to the observer of a hierarchical control system. The observer could tell that the system is controlling many variables simultaneously but he or she would find it difficult to demonstrate that some of these variables are being controlled in order to control others. For example, the observer could tell that a typist is controlling letter sequences, keypress events, finger movements and finger configurations. But the observer would have a hard time showing that these variables are hierarchically related. The observer could make up a plausible hierarchical description of these behaviors; for example, finger positions seem to be used to produce finger movements which are used to produce keypresses which are used to produce letter sequences. But finding a hierarchical description of behavior does not prove that the behavior is actually produced by a hierarchical process (Davis, 1976; Kline, 1983).

Hierarchical Invariance

Hierarchical production of behavior implies that the commands required to produce a lower level behavior are nested within the commands required to produce a higher level behavior. For example, the commands that produce a particular finger configuration would be nested within the commands that produce a movement from one configuration to another. Sternberg, Knoll and Turlock (1990) refer to this nesting as an invariance property of hierarchical control. Lower level commands are like a subprogram that is invoked by a program of higher level commands. The invariance of hierarchical control refers to the assumption that the course of such a subprogram does not depend on how it was invoked from the program (low level invariance); similarly, the course of the program does not depend on the nature of the commands carried out by the subprograms (high level invariance).

Convergent and Divergent Control. The hierarchical control model satisfies both the low and high level invariance properties of hierarchical control. The commands issued by higher level systems have no effect on the commands issued by lower level systems and vice versa. It is important to remember, however, that the commands in the control hierarchy are requests for input, not output. Higher level systems tell lower level systems what to perceive, not what to do. This aspect of control system operation solves a problem that is either ignored or glossed over in most hierarchical models of behavior: How does a high level command get turned into the the lower level commands that produce results that satisfy the high level command? If commands specify outputs then the result of the same command is always different due to varying environmental disturbances. The high level command to press a key, for example, cannot know which lower level outputs will produce this result on different occasions. This problem is solved by the hierarchical control model because intended results are represented as a convergent function rather than a divergent network.

Most hierarchical models of behavior require that a high level command be decomposed into the many lower level commands that produce the intended result. In the hierarchical control model, both the high level command and the intended result of the command are represented by a single, unidimensional signal. The signal that represents the intended result is a function of results produced by many lower level commands. But the high level command does not need to be decomposed into all the appropriate lower level commands (Powers, 1979). The difference between the high level command and the perceptual result of that command is sufficient to produce the lower level commands that keep the perceptual result at the commanded value (Marken, 1990).

Levels of Behavior

The hierarchical invariance properties of the control hierarchy provide a basis for determining whether its behavior is actually generated by hierarchical processes. Hierarchical control can be seen in the relative timing of control actions. In a control hierarchy, lower level systems must operate faster than higher level systems. Higher level systems cannot produce a complex perceptual result before the lower level systems have produced the component perceptions on which it depends. This nesting of control actions can be seen in the differential speed of operation of control systems at different levels of the control hierarchy. Lower level systems not only correct for disturbances faster than higher level ones; they carry out this correction process during the higher level correction process. The lower level control process is temporally nested within the higher level control process.

Arm Movement. Powers, Clark and McFarland (1960) describe a simple demonstration of nested control based on relative timing of control system operation. A subject holds one hand extended straight ahead while the experimenter maintains a light downward pressure on it. The subject is to move his or her arm downward as quickly as possible when the experimenter signals with a brief, downward push on the subject's extended hand. The result of this simple experiment is always the same: the subject responds to the downward signal push with a brief upward push followed by downward movement of the arm. An electromyograph shows that the initial upward push is an active response and not the result of muscle elasticity.

The arm movement demonstration reveals one level of control nested within another. The subject's initial upward push (which cannot be suppressed) is the fast response of a lower level control system that is maintaining the perception of arm position in a particular reference state (extended forward). The behavior of this system is nested within the response time of a higher level system that moves the arm downward. The higher level system operates by changing the reference for the arm position control system. The downward signal push causes the brief upward reaction because the signal is treated as a disturbance to arm position. This is particularly interesting because the signal is pushing the arm in the direction it should move; the lower level reaction is "counter productive" with respect to the goal of the higher level system (which wants to perceive the arm down at the side). The reaction occurs because the lower level system starts pushing against the disturbance to arm position before the higher level system can start changing the reference for this position.

Polarity Reversal. More precise tests of nested control were carried out in a series of experiments by Marken and Powers (1989). In one of these experiments, subjects performed a standard pursuit tracking task, using a mouse controller to keep a cursor aligned with a moving target. At intervals during the experiment the polarity of the connection between mouse and cursor movement was reversed in a way that did not disturb the cursor position. Mouse movements that had moved the cursor to the right now moved it to the left; mouse movements that had moved the cursor to the left now moved it to the right.

A sample of the behavior that occurs in the vicinity of a polarity reversal is shown in Figure 2. The upper traces show the behavior of a control system model and the lower traces show the behavior of a human subject. When the reversal occurs, both the model and the subject respond to error (the deviation of the cursor from the target) in the wrong direction, making it larger instead of smaller (any deviation of the error trace from the zero line represents an increase in error). The larger error leads to faster mouse movement which causes the error to increase still more rapidly. A runaway condition ensues with error increasing exponentially.

Figure 2 Here

About 1/2 second after the polarity reversal the subject's behavior departs abruptly from that of the model. The subject adjusts to the polarity reversal and the error returns to a small value. The model cannot alter its characteristics and the error trace quickly goes off the graph. These results provide evidence of two nested levels of control operating at different speeds. The faster, lower level system control the distance between cursor and target. This system continues to operate as usual even when, due to the polarity reversal, this causes an increase in perceptual error. Normal operation is restored only after a slower, higher level system has time to control the relationship between mouse and cursor movement.

Levels of Perception

The arm movement and polarity shift experiments reveal the hierarchical organization of control from the point of view of the observer. The hierarchical control model suggests that it should also be possible to view hierarchical organization from the point of view of the actor. From the actor's point of view, hierarchical control would be seen as a hierarchy of changing perceptions. One way to get a look at this hierarchy is again in terms of relative timing; in this case, however, in terms of the relative timing of the perceptual results of control actions rather of the actions themselves.

Computation Time Window. The hierarchical control model represents the results of control actions as unidimensional perceptual signals. A configuration, such as the letter "h", is a possible result of control actions, as is a sequence of letters, such as the word "hello". The model represents these results as perceptual input signals, the intensity of a signal being proportional to the degree to which a particular result is produced. This concept is consistent with the physiological work of Hubel and Wiesel (1979) who found that the firing rate of an afferent neuron is proportional to the degree to which a particular environmental event occurs in the "receptive field" of the neuron.

Many of the higher level classes of perception in the control hierarchy depend on environmental events that vary over time. Examples are transitions, events, and sequences. The neural signals that represent these variables must integrate several lower level perceptual signals that occur at different times. Hubel and Weisel found evidence of a computation time window for integrating perceptual signals. Certain cells respond maximally to configurations (such as "lines") that move across a particular area of the retina at a particular rate. These are "motion detector" neurons. The neuron responds maximally to movement of a configuration that occurs within a particular time window. Movement that occurs outside of this time window is not included in the computation of the perceptual signal that represents motion.

Levels by Time The hierarchical control model implies that the duration of the computation time window increases as you go up the hierarchy. The minimum computation time window for the perception of configurations should be shorter than the minimum computation time window for the perception of transitions which should be shorter than the minimum computation time window for the perception of sequences. I have developed a version of the psychophysical method of adjustment which makes it possible to see at least four distinct levels of perception by varying the rate at which items occur on a computer display. A computer program presents a sequence of numbers at two different positions on the display. The presentation positions are vertically adjacent and horizontally separated by 2 cm. The numbers are presented alternately to the two positions. The subject can adjust the rate at which the numbers occur in each position by varying the position of a mouse controller 1.

The results of this study are shown schematically in Figure 3. At the fastest rate of number presentation subjects report that the numbers appear to occur in two simultaneous streams. The fact that the numbers are presented to the two positions alternately is completely undetectable. However, even at the fastest rate of number presentation subjects can make out the individual numbers in each stream. At the fastest rate, there are approximately 20 numbers per second in each stream. This means that there is a 50 msec period available for detecting each number. This duration is apparently sufficient for number recognition suggesting that the computation time window for perception of configuration is less than 50 msec. Studies of the "span of apprehension" for sets of letters suggest that the duration of the computation time window for perception of visual configuration may be even less less than 50 msec, possibly as short as 15 msec (Sperling, 1960).

Figure 3 Here

As the rate of number presentation slows, the alternation between numbers in the two positions becomes apparent. Subjects report perception of alternation or movement between numbers in the two positions when the numbers in each stream are presented at the rate of about 7 per second. At this rate, an alternation from a number in one stream to a number in another occurs in 160 msec. This duration is sufficient for perception of the alternation as a transition or movement from one position to the other suggesting that the computation time window for transition perception is on the order of 160 msec. This duration is compatible with estimates of the time to experience optimal apparent motion when configurations are alternately presented in two different positions (Kolers, 1972).

The numbers presented in each stream are always changing. However, subjects find it impossible to perceive the order of the numbers as they alternate from one position to another even though it is possible to clearly perceive the individual numbers and the fact that they are alternating and changing across positions. The rate of number presentation must be slowed considerably, so that each stream of numbers is presented at the rate of about two per second, before it is possible to perceive the order in which the numbers occur. At this rate numbers in the sequence occur at the rate of four per second. These results suggest that the duration of computation time window for the perception of sequence is about 0.5 seconds. This is the time it takes for two elements of the sequence to occurQ the minimum number that can constitute a sequence.

The numbers in the rate adjustment study did not occur in a fixed, repeating sequence. Rather, they were generated by a set of rulesQ a program. The sequence of numbers was unpredictable unless the subject could perceive the rule underlying the sequence. The rule was as follows: if the number on the right was even then the number on the left was greater than 5, otherwise the number on the left was less than 5. (Numbers in the sequence were also constrained to be between 0 and 9). Subjects could not perceive the program underlying the sequence of numbers until the speed of the two streams of numbers was about .25 numbers per second so that the numbers in the program occurred once every two seconds. The perception of a program in a sequence of numbers requires considerably more time then it takes to perceive the order of numbers in the same sequence.

The perception of a sequence or a program seems to involve more mental effort than the perception of a configuration or a transition. Higher level perceptions, like programs, seem to represent subjective rather than objective aspects of external reality; they seem more like interpretations than representations. These higher level perceptions are typically called "cognitions". Of course, all perceptions represent subjective aspects of whatever is "out there"; from the point of view of the hierarchical control model, the location of the line separating perceptual from cognitive representations of reality is rather arbitrary. Behavior is the control of perceptions which range from the simple (intensities) to the complex (programs).

Perceptual Speed Limits.The hierarchical control model says that all perceptions of a particular type are controlled by systems at the same level in the hierarchy. This implies that the speed limit for a particular type of perception should be about the same for all perceptions of that type. The 160 msec computation time window for perception of transition, for example, should apply to both visual and auditory transition. There is evidence that supports this proposition. Miller & Heise (1950) studied the ability to perceive an auditory transition called a "trill". A trill is the perception of a temporal alternation from one sound sensation or configuration to another. The speed limit for trill perception is nearly the same as the speed limit for visual transition perception found in the number rate adjustment study -- about 15 per second. As in the visual case, when the rate of alternation of the elements of the auditory trill exceeds the computation time window the elements "break" into two simultaneous streams of sound; the perception of transition (trill) disappears even though the sounds continue to alternate.

There is also evidence that the four per second speed limit for sequence perception found in the number rate adjustment study applies across sensory modalities. Warren, Obusek, Farmer, & Warren (1969) studied subjects' ability to determine the order of the components sounds in a sound sequence. They found that subjects could not perceive the order of the components until the rate of presentation of the sequence was less than or equal to four per second. This was a surprising result because it is well known that people can discriminate sequences of sounds that occur at rates much faster than four per second. In words, for example, the duration of the typical phoneme is 80 msec so people can discriminate sequences of phoneme sounds that occur at the rate of about 10 phonemes per second. But there is reason to believe that the phonemes in a word are not heard as a sequence; that is, the order of the phonemes cannot be perceived. Warren (1974) showed that subjects can learn to tell the difference between sequences of unrelated sounds that occur at rates of 10 per second. However, the subjects could not report the order of the sounds in each sequence; only that one sound event differed from another. A word seems to be a lower order perception -- an event perception -- which is recognized on the basis of its overall sound pattern. There is no need to perceive the order in which the phonemes occur; just that the temporal pattern of phonemes (sound configurations) for one word differs from that for other words.

The Relationship Between Behavior and Perception

Configurations, transitions, events, sequences and programs are potentially controllable perceptions. An actor can produce a desired sequence of sounds, for example, by speaking sound events (phonemes) in some order. An observer will see the production of this sequence as a behavior of the actor. The hierarchical control model suggests that the actor's ability to produce this behavior turns on his or her ability to perceive the intended result. Since perception depends on speed, it should be impossible for the actor to produce an intended result faster than the result can be perceived. The observer will see this speed limit as a behavioral limit. An example of this can be seen in the arm movement experiment described above. In that experiment it appears that the time to respond to the signal push is a result of a behavioral speed limit; the inability to generate an output faster than a certain rate. But a closer look indicates that the neuromuscular "output" system is perfectly capable of responding to a signal push almost immediately, as evidenced by the immediate upward response to the downward signal push. The same muscles that produce this immediate reaction must wait to produce the perception of the arm moving downward. The speed limit is not in the muscles. It is in the results that the muscles are asked to produce; a static position of the arm (a configuration perception) or a movement of the arm in response to the signal push (a relationship perception).

Sequence Production and Perception. Some of the most interesting things people do involve the production of a sequence of behaviors. Some recent studies of temporal aspects of sequence production are directly relevant to the hierarchical control model. In one study, Rosenbaum (1989) asked subjects to speak the first letters of the alphabet as quickly as possible. When speed of letter production exceeded four per second the number of errors (producing letters out of sequence) increased dramatically, indicating a loss of control of the sequence. The speed limit for sequence production corresponds to the speed limit for sequence perception -- four per second.

The letter sequence study does not prove that the speed limit for letter sequence production is caused by the speed limit for letter sequence perception. It may be that the speed limit is imposed by characteristics of the vocal apparatus. However, in another study Rosenbaum (1987) found the same four per second speed limit for production of errorless finger tap sequences. The speed limit for finger tap sequence production is likely to be a perceptual rather than a motor limit because we know that people can produce finger taps at rates much higher than four per second. Pianists, for example, can do trills (alternating finger taps) at rates which are far faster than four per second. Further evidence of the perceptual basis of the finger tap sequence speed limit would be provided by studies of finger tap sequence perception. When a subject produces a sequence of finger taps he or she is producing a sequence of perceptions of pressure at the finger tips. A perceptual experiment where a pressure is applied to the tip of different fingers in sequence should show the four per second speed limit. Subjects should have difficulty identifying the order of finger tip pressures when the sequence occurs at a rate faster than four per second.

Confounding Levels. It is not always easy to find clear-cut cases of behavioral speed limits that correspond to equivalent perceptual speed limits. Most behavior involves the control of many levels of perception simultaneously. People control higher level perceptions (like sequences) while they are controlling lower level perceptions (like transitions). This can lead to problems when interpreting behavioral speed limits. For example, Rosenbaum (1983) presents some finger tapping results that seem to violate the four per second speed limit for sequence perception. When subjects tap with two hands they can produce a sequence of at least 8 finger taps per second. But each tap is not necessarily a separate event in a sequence. Some pairs of taps seem to occur at the rate at which sequences are experienced as events. A sequence of finger taps is an event in the same sense that the sequence of muscle tensions that produce a finger tap is an event; the order of the components of the sequence cannot be perceived. These finger tap events are then unitary components of the sequence of finger tap perceptions.

The fact that certain pairs of finger taps are produced as events rather than ordered sequences is suggested by the errors made at each point in the finger tap sequence. Errors occur most frequently at the point in the sequence at which a fast pair is being initiated. Errors rarely occur for the second element of a fast pair. This suggests that the errors occur at the sequence level rather than the event level. The subject's attempts to produce a keypress sequence too rapidly apparently interferes with sequence rather than event production. Events are already produced at a fast enough rate and an increase in the speed of sequence production has little effect on the ability to control the component events.

Changing Perception Can Change Behavior:SGoing Up A LevelS. The relationship between perception and behavior can be seen when a person learns to perform a task by controlling a new perceptual variable. An example of this can be seen in simple pursuit tracking tasks. In the typical tracking task the target moves randomly. When, however, a segment of target movement is repeated regularly the subject's tracking performance improves markedly with respect to that segment (Pew, 1966). According to the hierarchical control model, control is improved because the repeated segment of target movement can be perceived as a predictable event. With the random target the subject must wait to determine target position at each instant in order to keep the cursor on target. With the repeated target, the subject controls at a higher level. keeping a cursor movement event matching a target movement event. The fact that the subject is now controlling a higher level perception (an event rather than a configuration) is evidenced by the longer reaction time when responding to a change in target movement. When controlling the target-cursor configuration the subject responds almost immediately to changes in target position. When controlling target-cursor movement events it takes nearly 1/2 second to respond to a change to the same change in target movement pattern.

An experiment by Robertson and Glines (1985) also shows improved performance resulting from changed perception. Subjects in the Robertson and Glines study performed a learning task where the solution to a computerized game could be perceived at several different levels. Subjects who were able to solve the game showed three distinct plateaus in their performance. The level of performance, as indicated by reaction time measurements, improved at each succeeding plateau. Because the same outputs (keypresses) were produced at each level of performance, each performance plateau was taken as evidence that the subject was controlling a different perceptual variable.

Behavior/Perception Correlations. Few psychologists would be surprised by the main contention of this paper: that there is an intimate relationship between perception and behavior. However, most models of behavior assume that the nature of this relationship is causal: behavior is guided by perception. This causal model provides no reason to expect a relationship between the structure of perception and behavior: no more than there is to expect a relationship between the structure of computer input and output. This does not mean that there might not be such a relationship; it is just not demanded by the causal model.

The control model integrates perception and behavior with a vengeance. Behavior is no longer an output but, rather, a perceptual input created by the combined effects of the actor and the environment. Behavior is perception in action. From this point of view, behavioral skills are perceptual skills. Thus, it is not surprising to find some indication of a correlation between behavioral and perceptual ability. For example, Keele and his colleagues (Keele, Pokorny, Corcos and Ivry, 1985) have found that the ability to produce regular time intervals between actions is correlated with with ability to perceive these intervals. These correlations were fairly low by control theory standards but they are expected if the production of regular time intervals involves control of the perception of these intervals.

Conclusion

This report has presented evidence that human behavior involves control of a hierarchy of perceptual variables. There is evidence that the behavior of non-human agents, such as chimpanzees, also involves the control of a similar hierarchy of perceptions (Plooij and van de Rijt-Plooij, 1990). A model of hierarchical control shows how studies of perception and behavior provide evidence about the nature of control from two different perspectives. Perceptual studies provide information about the ability to perceive potentially controllable consequences of actions. Behavioral studies provide information about the ability to produce desired consequences. The factors that influence the ability to perceive the consequences of action should also influence the ability to produce them. In both cases we learn something about how agents control their own perceptions.

The hierarchical control model shows that limitations on the ability to produce behavior may reflect limitations on the ability to perceive intended results. The speed at which a person can produce an errorless sequence of events, for example, is limited by the speed at which the order of these events can be perceived. But not all skill limitations are perceptual limitations. Controlled (perceived) results are produced, in part, by the outputs of the behaving agent. The ability to produce certain outputs can limit the ability to control certain perceptions. For example, it is imposs- ible to perceive oneself lifting a 300 pound barbell until the muscles have been developed to the point that they are able to generate the output forces necessary to control this perception.

Perception and behavior are typically treated as two completely different types of phenomena. Perception is a sensory phenomenon: behavior is a physical phenomenon. But the concept of control as the behavior of perception suggests that this separation is artificial. Perception and behavior are the same phenomenon seen from two different perspectives. In order to understand how this phenomenon works, it will be necessary to understand how agents perceive (perception) and how they act to affect their perceptions (behavior). Studies of perception and behavior should become an integral part of the study of a single phenomenonQ control.

References

Albus, J. Brains, behavior and robotics. Petersborough, NH: Byte Books

Arbib, M. (1972) The metaphorical brain. New York: Wiley

Bryan, W. L. and Harter, N. (1899) Studies on the telegraphic language: The acquisition of a hierarchy of habits. Psychological Review, 6, 345-375

Davis, W. J. (1976) Organizational concepts in the central motor network of invertebrates. In R. M. Herman, S. Grillner, P.S.G. Stein, & D. Stuart (Eds.) Neural control of locomotion, New York: Plenum (p. 265)

Greeno, J.G. and Simon, H.A. (1974) Processes for sequence production. Psychological Review, 81, 187-197

Herrigal, E. (1971) Zen in the art of archery. New York: Vintage

Hubel, D. H. and Wiesel, T. N. (1979) Brain mechanisms of vision. In J.M. Wolfe (Ed.) The mind's eye: Readings from Scientific American, New York: Freeman (pp. 40-52)

Keele,S.W., Cohen, A. & Ivry, R. I. (1990) Motor programs: Concepts and issues. In M. Jeannerod (Ed.) Attention and performance XIII: Motor representation and control. New Jersey: Erlbaum. (pp. 77-110)

Keele,S.W., Pokorny, R.A.,Corcos,D.M. & Ivry, R. I. (1985) Do perception and production share common timing mechanisms: A correlational analysis. Acta Psychologica, 60, 173-191

Kline, R. (1983) Comment on Rosenbaum et al. Hierarchical control of rapid movement sequence. Journal of Experimental Psychology: Human Perception and Performance, 9, 834-36. ?1;2c Kolers, P.(1972) The illusion of movement. In R. Held and W. Richards (Eds) Perception: Mechanisms and models, San Francisco:Freeman.

Lashley, K.S. (1951) The problem of serial order in behavior. In L.A. Jeffress (Ed.) Cerebral mechanisms in behavior. New York: Wiley

Marken, R.S. (1986) Perceptual organization of behavior: A hierarchical control model of coordinated action. Journal of Experimental Psychology: Human Perception and Performance, 12, 267-76.

Marken, R. S. (1988) The nature of behavior: Control as fact and theory, Behavioral Science, 33, 196-206

Marken, R. S. (1989) Behavior in the first degree.In W. Hershberger (Ed) Volitional Action: Conation and control, Amsterdam: North- Holland

Marken, R. S. (1990) Spreadsheet analysis of a hierarchical control system model of behavior, Behavior Research Methods, Instruments, & Computers, 22, 349 - 359.

Marken, R. S. and Powers, W. T. (1989) Levels of intention in behavior. In W. Hershberger (Ed.) Volitional Action: Conation and Control, Elsevier Science Publishers: North-Holland.

Martin, J. G. (1972) Rhythmic (hierarchical) versus serial structure in speech and other behavior. Psychological Review, 79, 487-509

Miller, G. A. and Heise, G. A. (1950) The trill threshold, Journal of the Acoustical Society of America, 22, 637-638

Palmer, S. E. (1977) Hierarchical structure in perceptual representation. Cognitive Psychology, 9, 441-474

Pew, R. W. (1966) Acquisition of hierarchical control over the temporal organization of a skill. Journal of Experimental Psychology, 71, 764-771

Plooij, F.X. and van de Rijt-Plooij, H.H.C.(1990) Developmental transitions as successive reorganizations of a control hierarchy. American Behavioral Scientist, 34, 67-80

Povel, D-J (1981) Internal representation of simple temporal patterns. Journal of Experimental Psychology: Human Perception and Performance, 7, 3 - 18

Powers, W. T. , Clark, R. K. and McFarland, R. L. (1960) A general feedback theory of human behavior: Part II. Perceptual and Motor Skills, 11, 309-323

Powers, W. T. (1973) Behavior: The control of perception. New York: de Gruyter

Powers, W. T. (1979) The nature of robots: Part 4: A closer look at human behavior. Byte, August, 309-323

Powers, W. T. (1989) Living control systems. Gravel Switch, KY: CSG Publishing

Powers, W. T. (1990) A hierarchy of control. In R. J. Robertson and W. T. Powers (Eds.) Introduction to modern psychology: The control- theory view, CSG Publishing: Gravel Switch, Ky.

Robertson, R. J. and Glines, L.A. (1985) The phantom plateau returns. Perceptual and Motor Skills, 61, 55-64

Robertson, R. J. and Marken, R. S. (1990) Perception: Input to the control function. In R. J. Robertson and W. T. Powers (Eds.) Introduction to modern psychology: The control-theory view, CSG Publishing: Gravel Switch, Ky.

Rosenbaum, D.A. (1987) Hierarchical organization of motor programs. In S.P.Wise (Ed) Higher brain functions: Recent explorations of the brain's emergent properties, New York:Wiley

Rosenbaum, D.A., Weber, R.J. Hazelett, W. M and Hundorff, V. (1986) The parameter remapping effect in human performance: Evidence from tongue twisters and finger fumblers. Journal of Memory and Language, 25, 710-725

Sperling, G. (1960) The information available in brief visual presentations. Psychological Monographs, 74 (Whole No. 11)

Sternberg, S., R.L. Knoll, S. Monsell, and C.E. Wright (1989) Motor programs and hierarchical organization in the control of rapid speech. In M. Jeannerod (Ed.) Attention and performance XIII: Motor representation and control. New Jersey: Erlbaum. (pp. 3-55)

Sternberg, S., R.L. Knoll, D.L. Turock (1990) Hierarchical control in the execution of action sequences. Phonetica, 45, 175 - 197

Warren, R. M. (1974) Auditory temporal discrimination by trained listeners, Cognitive Psychology, 6, 237-256

Warren, R.M., Obusek, C.J., Farmer, R.M. and Warren, R.P. (1969) Auditory sequence: Confusion of patterns other than speech or music. Science, 164, 586-587

Footnote

1. A HyperCard version of the number rate adjustment program is available from the author. The program will be sent upon receipt of a formatted 3 1/2- in. double density or high density disk in a reusable mailer with return postage.

Figure Captions

Figure 1. Perceptual Control Hierarchy (after Powers, 1989, p 278)

Figure 2. Lower level runaway response to mouse-cursor polarity reversal.

Figure 3. Schematic representation of the results of the number rate adjustment study.

Date: Fri Oct 16, 1992 11:07 am PST
Subject: Re: Apparent S-R behaviour: what goes on?

Perhaps there is. The response in a situation where one needs to say "Sorry" should be immediate, otherwise there is suspicion that the response was not representative of real feelings. Shannon, and for that matter Wiener remarked essentially that the longer you can accumulate channel traffic before trying to understand it, the better your rejection of noise is going to be. The need for rapid response always ups the error rate. Hence the Iranian Airbus. A central question is whether to push your hypothesis tester towards type A or type B errors, i.e. in a street gunfight to risk being shot or to risk shooting an innocent bystander.

John Gabriel (gabriel@eid.anl.gov)

Date: Fri Oct 16, 1992 12:05 pm PST
Subject: CSGnet Growth

[from Gary Cziko 921016.1800 GMT]

Since there have been several new subscribers to CSGnet over the last few days (including Doktor Delirium from Brazil), I thought I would take a look at the latest list (this can be gotten by anyone by sending the command REV CSG-L (COUNTRIES as the first line of a message to LISTSERV@VMD.CSO.UIUC.EDU (please do not send this command to CSGnet).

To my surprise, I discovered that the list of subscribers is now conveniently organized into countries. I've appended a summary of how many and where CSGnetters are located. This does NOT include any individuals who read and/or interact with CSGnet via Usenet (NetNews)--I have no way of knowing how many or where these people are.

So as you can see, the sun never sets on CSGnet. And we continue to grow at a healthy rate.--Gary

```
*  Country        Subscribers
*  -------        -----------
*  Australia          1
*  Belgium            1
*  Brazil             1
*  Canada            13
*  Chile              1
*  France             3
*  Germany            4
*  Great Britain      3
*  Israel             2
*  Mexico             1
*  Netherlands        4
*  Poland             1
*  Singapore          1
*  Switzerland        1
*  Taiwan             1
*  USA               87
*  Yugoslavia         1
*
* Total number of users subscribed to the list:  126
* Total number of countries represented:          17
```

Date: Fri Oct 16, 1992 12:49 pm PST
Subject: Modeling reorganization

[From Bill Powers (921016.0930)]   Oded Maler (921016) --

I like it. We develop a verbal means of protecting ourselves from criticism or resistance in bumping our way among people, and use it automatically when we bump into anything, or cause any problem. I've sometimes thought that "pardon me" is used as another way of saying "get out of my way."

What kind of experiment could you do to test your proposal that the "sorry" loop is lower in the hierarchy than the conscious one?

Any more proposals out there? How about the S-R explanation?

Greg Williams (921016) --

>... the criteria for starting reorganization due to errors in
>(some? all?) innate critical variables are subject to alteration by
>reorganization triggered by sufficient error in acquired (within a
>single lifetime) reference levels -- which can be thought of as
>acquired critical variables.

Do I understand this correctly? You appear to be saying that a learned system can alter not just the states of the organism that I call critical variables (through indirect effects of its actions), but the TARGET VALUES (criteria) toward which those critical variables are controlled and that define zero critical error. So unless you're junking my proposed reorganizing system, you're adding lines in the diagram that come from the output of a learned reorganizing system and go to the reference inputs of my proposed system. Is that what you mean?

I know you're trying to get the house finished and do all the other things that would keep three ordinary people busy. When things settle down, perhaps you can get more specific about the model you're proposing. I don't really want to go on with this until there's a model to test. What you consider to be data depends on the model you're using.

Best, Bill P.

Date: Fri Oct 16, 1992 1:04 pm PST
Subject: workshop

U S C

Neural Mechanisms of Looking, Reaching and Grasping

Workshop sponsored by the Human Frontier Science Research Program and the Center for Neural Engineering - U.S.C.

Michael A. Arbib Organizer

October 21-22, 1992 HEDCO NEUROSCIENCES AUDITORIUM
USC, University Park Campus, Los Angeles, CA

Session 1, October 21 Chair: Hideo Sakata

08:30 - 09:00 am Marc Jeannerod (INSERM, Lyon, France)

09:00 - 09:30 am "Functional Parcellation of Human Parietal and Premotor Cortex during Reach and Grasp Tasks" Scott Grafton School of Medicine, USC, Los Angeles, CA, USA

09:30 - 10:00 am "Anatomo-functional Organization of the 'Supplementary Motor Area' and the Adjacent Cingulate Motor Areas" Massimo Matelli Universita Degli Studi di Parma, Italy

10:30 - 11:00 am "Inferior Area 6: New findings on Visual Information Coding for Reaching and Grasping" Giacomo Rizzolatti, Universita Degli Studi di Parma, Italy

11:00 - 11:30 am "Neural Strategies for Controlling Fast Movements" Jim-Shih Liaw CNE/Computer Science Department, USC Los Angeles, CA, USA

11:30 - 12:00 am "Cortex and Haptic Memory" Joaquin Fuster, UCLA Medical Center Los Angeles, CA, USA

12:00 - 12:30 pm   "Trajectory Learning from Spatial Constraints" Michael Jordan Brain and Cognitive Science Department MIT, Cambridge, MA, USA

Chair: Jean-Paul Joseph

01:30 - 02:00 pm   "Selectivity of Hand Movement-Related Neurons of the Parietal Cortex in Shape, Size and Orientation of Objects and Hand Grips" Hideo Sakata, Nihon University School of Medicine Tokyo, Japan

02:00 - 02:30 pm   "Modeling the Dynamic Interactions between Subregions of the Posterior Parietal and Premotor Cortices" Andrew Fagg CNE/Computer Science Department, USC, Los Angeles, CA, USA

02:30 - 03:00 pm   "Optimal Control of Reaching Movements Using Neural Networks" Alberto Borghese Center for Neural Engineering, USC and I.F.C.N.-C.N.R., Milano, Italy

**** 03:00 - 03:30   BREAK

03:30 - 04:00 pm   " How the Frontal Eye Field can impose a saccade goal on Superior Colliculus Neurons" Madeleine Schlag-Rey Brain Research Institute, UCLA, Los Angeles, CA, USA

04:00 - 04:30 pm   "Variations on a Theme of Hallett and Lightsone" John Schlag Department of Anatomy, UCLA Los Angeles, CA, USA

04:30 - 05:00 pm   "The saccade and its Context" Lucia Simo Center for Neural Engineering, USC, Los Angeles, CA

05:00 - 05:30   "An Integrative View on Modeling" Michael Arbib Center for Neural Engineering/Computer Science Department, USC Los Angeles, CA, USA

Chair: Giacomo Rizzolatti

08:30 - 09:00 am "Neural Activity in the Caudate Nucleus of Monkeys during Motor and Oculomotor Sequencing" Jean-Paul Joseph INSERM, Lyon, France

09:00 - 09:30 "Models of Cortico-Striatal Plasticity for Learning Associations in Space and Time" Peter Dominey Computer Science Department, USC, Los Angeles, CA, USA

09:30 - 10:00 "Eye-Head-Hand Coordination in a Pointing Task" Claude Prablanc INSERM, Lyon, France

10:30 - 11:00 "Modeling Kinematics and Interaction of Reach and Grasp" Bruce Hoff CNE/Computer Science Department, USC, Los Angeles, CA, USA

11:00 - 11:30 "Towards a Model of the Cerebellum" Nicolas Schweighofer Center for Neural Engineering, USC, Los Angeles, CA, USA

11:30 - 12:00 "Does the Lateral Cerebellum Map Movements onto Spatial Targets?", Thomas Thach Washington University School of Medicine, St. Louis, MO, USA

Date: Sat Oct 17, 1992 2:43 am PST
Subject: Time out

From Greg Williams (921017)   >Bill Powers (921016.0930)

>You appear to be saying that a learned system can alter not just the states
>of the organism that I call critical variables (through indirect effects of
>its actions), but the TARGET VALUES (criteria) toward which those critical
>variables are controlled and that define zero critical error. So unless
>you're junking my proposed reorganizing system, you're adding lines in the
>diagram that come from the output of a learned reorganizing system and go
>to the reference inputs of my proposed system. Is that what you mean?

I think so. I'm not saying I've got it all mapped out in detail, but I'm saying it looks like something of this sort is needed to account for situations such as someone "gladly" killing oneself "for one's country," (some) "heroes" getting into (perceived by them as) perilous straits to save other persons, and someone "choosing death before dishonor." Still, this point isn't on the mainline of our argument, but a sidetrack which I don't judge of utmost importance.

>I know you're trying to get the house finished and do all the other things
>that would keep three ordinary people busy. When things settle down, perhaps
>you can get more specific about the model you're proposing. I don't really
>want to go on with this until there's a model to test. What you consider to
>be data depends on the model you're using.

Sounds reasonable; I agree. I also think that the models one builds must be fully informed by genuine data on what is being modeled, in addition to meeting criteria of internal consistency, elegance, and being informed by data on what might or might not be analogous to what is being modeled. Otherwise, the modeler could end up with absurd claims like that made by the (apocryphal?) aerodynamicist whose model said that bumblebees can't fly. So, while I attend to other business for a while, I suggest that both of us take data on human education seriously: some teachers appear to aid the learning of some students in some subject areas. Both of us should be able to show how our models explain that data, regardless of what they explain about data on E. coli.

Best wishes,   Greg

Date: Sat Oct 17, 1992 3:55 am PST
Subject: "sorry"

This is just an example, not a theory:

On at least a few occasions in my life, I've bumped my head or stubbed my toe and said, "Sorry!" or the equivalent.

Best   Eileen

Date: Sat Oct 17, 1992 7:53 am PST

[From Bill Powers (921017.0900)]   Greg Williams (921017) --

>>you're adding lines in the diagram that come from the output of a
>>learned reorganizing system and go to the reference inputs of my
>>proposed system. Is that what you mean?

>I think so. I'm not saying I've got it all mapped out in detail, but I'm
>saying it looks like something of this sort is needed to account for
>situations such as someone "gladly" killing oneself "for one's country,"
>(some) "heroes" getting into (perceived by them as) perilous straits to
>save other persons, and someone "choosing death before dishonor." Still,
>this point isn't on the mainline of our argument, but a sidetrack which
>I don't judge of utmost importance.

This "sidetrack" may be more important than you think. It indicates to me that you haven't understood my proposal about a reorganizing system -- that what you're arguing against isn't even what I'm proposing.

My reorganizing system has no reference levels concerning "life" or "death," or "peril" or "danger" or "survival." Those concepts belong in the learned systems. I can understand that there can be conflict between learned goals, such as "safety" and "patriotism," and that the control system with the highest gain and output capacity will win the conflict (if any side wins it). But these goals do not conflict with critical reference levels. They themselves have nothing to do with critical reference levels. They are learned cognitive goals in the upper levels of the hierarchy.

Critical variables are related only to actual present-time states of the organism itself. The critical reference signals likewise have nothing specific to do with events external to the organism or relations of the organism to external events or things. The reorganizing system has no fear of death or hope for survival. It recognizes neither danger nor safety. It does not even know that food is a good thing to eat, or that water is a good thing to drink. It is concerned strictly with the current state of the organism itself, in terms that can have meaning before any organized hierarchy exists, before the organism has any perceptions of an "outside world" or a "body." It will just as readily put the organism in danger as save it from danger, because it knows nothing of danger in the world outside the organism.

Consider the soldier going into combat. This soldier wishes to defend his country, and he also wishes to stay alive. Both of these goals are learned; they are cognitive goals. The soldier has also learned that combat could end his life, that bravery is something to be sought, and that desertion in the face of the enemy is severely punished, that valor is copiously rewarded. So the soldier wants to go into combat and he wants not to go into combat. He feels a desire to flee, and a desire to stay, with the result that he experiences fear and other emotions. If the degree of conflict is sufficient to cause critical error, the soldier will begin reorganizing. He may resolve the conflict by learning to perceive an honorable death as glorious (like Worf), or to consider it non-threatening, or to imagine that something protects him from death, or to believe that he will simply awaken into a better world after death. That would remove the conflict and the critical error caused by being in conflict. The soldier would then march to his death without disturbing any critical variables. Of course the outcome might go the other way; the soldier might reorganize so that the concept of patriotism is modified or abolished, leaving the self-preservation side in charge. He will then gladly desert and suffer the symbolic punishment, and again the critical error will be corrected.

You can see that I would wonder why you consider it necessary for a learned system to change a critical reference signal.

>I also think that the models one builds must be fully informed by genuine
>data on what is being modeled, in addition to meeting criteria of internal
>consistency, elegance, and being informed by data on what might or might
>not be analogous to what is being modeled.

What you consider "genuine data" depends on the model you already believe in. If you come into the discussion thinking that the outside world can "facilitate" control, then you will interpret the relationship of the outside world to the organism as demonstrating "facilitation," something done to the organism in a purposeful way. If, on the other hand, you think that all changes are internally motivated and accomplished, you will see the very same actions in the outside world as having a different meaning in relation to the organism. You will see the "facilitator" as doing nothing more than applying disturbances and rearranging the environment, without any special effect on the organism. You will see all changes in behavior as resulting from the natural adjustments of the organism to disturbances and changes in feedback parameters, with the causes of these changes being entirely internal, beyond influence by anything external.

This is similar to the difference in the way S-R psychologists and CTers see "reinforcement." To the S-R psychologist, a reinforcer has some effect on the organism that alters the way it responds to stimuli. The underlying model is causal; reinforcers have a special kind of influence on organisms. So the S-R psychologist can produce mountains of "genuine data" showing how reinforcers have these effects on the responses of organisms (defined by their outcomes). The data seem to support the concept of reinforcement because they are interpreted in exactly the way needed to make them seem to do that.

The CTer, of course, sees reinforcements as controlled variables. When the CT theorist looks at the same genuine data, it does not at all support the idea of a special kind of effect on behavior. Now the data simply show that the organism produces whatever behavior is required to keep the so-called reinforcer at a reference level determined by the organism. The reinforcer is no longer seen as having any special effect on the organism, other than the effects that any sensory input has.

Data are "genuine" only when described at a sufficiently low level of abstraction that there is no disagreement on what is being observed -- where theory doesn't enter into the description. To say anything more than what was physically done is to bring in theories. So I can say that I present an organism with a certain environmental situation. I cannot say, without aid of a theory, that doing this amounted to "facilitation" or "teaching" or any of those abstract terms that carry within them a theory of causality, that imply an arrow reaching from the outside world into the organism and changing something inside the organism. All descriptions that imply an effect inside an organism controlled by an external agency are based on the S-R concept of behavior and the causal model that supports it. All data concerning such effects, however genuine, can also be interpreted a different way under control theory.

Best, Bill P.

Date: Sat Oct 17, 1992 8:18 am PST
Subject: Delprato to Kurtzer re. Skinner

[From: Dennis Delprato] (isaac kurtzer)

> reading your article am i to understand that skinner felt that
> psychology could be reduced to chemistry,physics,etc. but that it was not
> necessary to explain behavior (superfluous?) ?

First, note the data reveal Skinner to be inconsistent on reductionism-nonreductionism. This is clearly discussed in the article.

On explanation: His view is that if one can predict and produce (control) a phenomenon, then this is explanation. There is nothing left to explain. One has done everything necessary. It is when one cannot predict and control events that we find the sort of theorizing he deplored (see discussion under "Purpose of Science."

When he uses the expressions "understanding," "understood," and "knowledge" (in several included quotations) he is addressing explanation. Furthermore, he uses the term "explanation" in the quotation on p. 2, col. 1, para. 2 (from his 1950, "Are Theories of Learning Necessary?", p. 193). What kind of explanation/theory did he find impedes science? Those with nonspatiotemporal referents/mentalistic ones/ spookological ones/supernatural ones.

Dennis Delprato Psy_Delprato@emunix.emich.edu

Date: Sat Oct 17, 1992 12:21 pm PST
Subject: Progress report, Misc.

[From Dag Forssell (921017-1)]

It has been over a week now since Christine and I prepared and presented the third day of our seminar. It, too went rather well. The engineers continued to be attentive, and several "got it."

We learned a lot from this opportunity to present our programs and even got paid a token amount. There is lots of room for improvement, and we are working on it. At the moment, I am continuing work on presenting the Behavior of Perception.

Also polishing details on demodisk programs. The documents that go with Bill Powers' programs have been edited to reflect current address. I am writing a description for the E Coli program and tumble, and am pulling together the documentation for Rick Markens spreadsheet program, with Rick's encouragement. I will make the demodisk available when I am done with it.

Some of you will remember Toto Grandes del Mazo from a conference four or more years ago. I met him in Jan of 1990 at a Deming seminar. He told me he had been trying to tell Dr. Deming about Control Theory (in vain) for a long time. He now lives in Los Angeles and consults somehow with the people who put on the Deming seminars. He also "assists" at the seminars, whatever that is. Last Monday, I met him again. He lectured a month ago and again this month to a "Deming users group" in Los Angeles on psychology. Before the lecture, they showed a video of last month's lecture, which we were unable to attend. In the video, he stated that Control Theory guides his thinking and specifically mentioned a William T. Powers as the source. He is actually quite good, asking his audience for their reference perceptions first, so he can address them, then at the end showing that he followed the control model in doing so.

During break, we connected and I brought him personal greetings from Bill. Later he introduced me to the audience as another teacher of Control Theory. Now, they want to have me as a lecturer. I had actually given up on marketing myself through that audience. It is funny how the ball bounces.

Gary has asked me to review the "starter document." To check it out, I have learned more about Bill Silvert's listserver, and how to access the files there. An updated document and some files will be mailed to Gary and Bill by Monday.

Rick's file on the spreadsheet (csg/marken/marken.doc on Bill's server) has no pictures. Therefore, I have just drawn them in ASCII. Here they are for any suggestions on improvements by Rick or anyone else.

____

FIGURE 1: A BASIC CONTROL SYSTEM

```
                              r
                              |
                        ______v_____
                       |            |
                  o-->| comparator |---o
                p |   |____________|   | e
             _____|_____          _____v_____
            |           |        |           |
            |  sensor   |        | amplifier |        System
_____________|___________|________|___________|_________________
                  ^                    |
                  |                    v          Environment
          d------> i <---------------- o


      |  = down        ^   = up        o-->  = connection and
      v    arrow       |     arrow     |       side arrow
_________________________________________________________________
```

____

FIGURE 2: A THREE LEVEL HIERARCHY OF CONTROL SYSTEMS

Only the left half of figure 2 shown here. The right half extends the picture to the full four "stacks" of interconnected, interrelated control systems.

```
                    r                                   r
                    |                                   |
Level 3       ______v_____                        _____v______
             |            |                      |            |
         o-->| comparator |---o              o-->| comparator |---o
       p |   |____________|   | e          p |   |____________|   | e
    _____|_____          _____v_____    _____|_____          _____v_____
   |           |        |           |  |           |        |           |
   |  sensor   |        | amplifier |  |  sensor   |        | amplifier |
   |___________|        |___________|  |___________|        |___________|
         ^                    |              ^                    |
         |                    |              |                    |
         o------------------- | -------------o------------------- | --------
         |                    |              |                    |
         |          o---------o------------- | --------o----------o---------
         |          | r                      |         | r
Level 2  |    ______v_____                   |    _____v______
         |   |            |                  |   |            |
         o-->| comparator |---o              o-->| comparator |---o
       p |   |____________|   | e          p |   |____________|   | e
    _____|_____          _____v_____    _____|_____          _____v_____
   |           |        |           |  |           |        |           |
   |  sensor   |        | amplifier |  |  sensor   |        | amplifier |
   |___________|        |___________|  |___________|        |___________|
```

```
         ^                    |               ^                    |
         |                    |               |                    |
         o------------------- | --------------o------------------- | --------
         |                    |               |                    |
         |          o---------o-------------- | --------o----------o---------
         |          | r                       |         | r
Level 1  |    ______v_____                    |    _____v______
         |   |            |                   |   |            |
         o-->| comparator |---o               o-->| comparator |---o
       p |   |____________|   | e           p |   |____________|   | e
    _____|_____          _____v_____     _____|_____          _____v_____
   |           |        |           |   |           |        |           |
   |  sensor   |        | amplifier |   |  sensor   |        | amplifier |
___|___________|________|___________|___|___________|________|___________|___
         ^                    |               ^                    |
         |                    v               |                    v
d------> i <----------------- o    d--------> i <----------------- o    d----

Legend:              ^                                 |
                     |                            ---- | ---
              o------o--->                             |
              |      |                                 v
           Connections, arrows                  Crossing paths,
              up & sideways                       arrow down
```

FIGURE 3. THE BASIC CONTROL SYSTEM AS IMPLEMENTED IN THE SPREADSHEET:

```
            Cell         Cell
            Values       Formulas
            ________
   R(j,i)  |      7 |     +E4-D4
   P(j,i)  |  6,004 |     @INDEX(PW,1,PW2)*P11+@INDEX(PW,2,PW21)*P12...
   O(j,i)  |  41.63 |     +O21+SLOW*(GAIN*(R21-P21)-O21)
            --------
```

FIGURE 4. PERCEPTUAL WEIGHTING MATRIX, PW:

PW Matrix

	1	1	1	1	1
	2	1	1	1	-1
	3	1	1	-1	1
	4	1	1	-1	-1
	5	1	-1	1	1
	6	1	-1	1	-1
	7	1	-1	-1	1
Row	8	1	-1	-1	-1
Labels	9	-1	1	1	1
	10	-1	1	1	-1
	11	-1	1	-1	1
	12	-1	1	-1	-1
	13	-1	-1	1	1
	14	-1	-1	1	-1
	15	-1	-1	-1	1
	16	-1	-1	-1	-1

FIGURE 5.

	A	B	C	D	E	F	G	H	I
1			System (i)		1	2	3	4	Average error
2									
3		Delay	Gain	R(3,i)	-1.00	1.00	1.00	1.00	
4	Level 3			P(3,i)	-1.00	1.00	1.00	1.00	0.000
5		1E-05	1000	O(3,i)	-11.60	38.58	58.58	74.21	
6									
7				R(2,i)	-11.60	50.19	20.00	15.63	
8	Level 2			P(2,i)	-11.80	49.97	20.17	15.63	0.140
9		0.0001	500	O(2,i)	21.86	17.98	0.53	21.25	
10									
11				R(1,i)	3.88	19.12	17.90	-16.80	
12	Level 1			P(1,i)	3.87	19.10	17.90	-16.80	0.010
13		0.01	50	O(1,i)	53.87	-69.10	67.89	-33.20	
14	System								
15	Input Variable:			I	3.87	19.11	17.89	-16.80	0.050
16									
17	Disturbance:			D	-50.00	50.00	-50.00	50.00	
18	Test variable: 12.00			Weights	-1.00	1.00	-1.00	1.00	
19	Stability								
20	Factor: 327.2			Behavior		19.49			

Now, I just have to learn to run the program too, and all will be well.

Best to all,   Dag

Date: Sun Oct 18, 1992 3:08 am PST
ubject: Finis?

From Greg Williams (921018)   >Bill Powers (921017.0900)

>You can see that I would wonder why you consider it necessary for a
>learned system to change a critical reference signal.

Yes, I can see. What makes most sense to me now, in light of your critique, is to hypothesize that the error criteria (to start reorganization) for some or all innate critical reference signals can be altered by learned systems. That would explain the patriot steadfastly refusing to tell state secrets under torture, I think; the patriot's learned "honor" (or whatever) system increases the error criterion for pain sufficiently that reorganization is not triggered up to the point of death. But maybe you think "not experiencing pain" is NOT an innate critical reference signal? In that case, I suppose I'd have to postulate that learned systems can alter the connections between perceptions such as "pain" and the innate critical reference systems which those perceptions (or higher-level combinations of them) reach. In this hypothesis, some feedback loops of innate critical reference signals could be disabled, in effect, by a learned system "breaking" the loop (i.e., in a torture situation, the "honor" system "turns off" the link between pain and some critical variables; of course, in other situations, the link is not "turned off").

Perhaps you should list what you count as innate critical variables, and then say which of these you think are triggering reorganization (if they do) in (1) the recent "learning a new phone number" example, (2) the "Fire!" example, and (3) my "tortured patriot" example.

I suspect -- on the basis of observations, rather than on modeling considerations (though efficiency considerations also might be invoked in this respect) -- that non-trial-and-error processes are important in at least SOME kinds of learning /reorganization (for example, in "being taught to swim"). "Overriding" of innate critical variables is just one way in which this can occur, but it isn't the only way (I claim). That's why I say "overriding" is a side-issue in what I consider to be our main argument about the extent to which the outcomes of a person's learnings/reorganizations depend on the person's environment (living/nonliving). I certainly could be wrong about the need for some acquired-in-a-lifetime systems to be capable of "overriding" innate critical variables; but if I am, there can still be environmental effects on the outcome of reorganization.

>What you consider "genuine data" depends on the model you already believe in.

Certainly. I was trying to point out the dangers in using data on E. coli behavior, and NOT using data on human social interactions, when attempting to build and test models of human social interactions. How can we test our models of reorganization with PEOPLE, rather than with bacteria?

>If, on the other hand, you think that all changes are internally motivated
>and accomplished, you will see the very same actions in the outside world as
>having a different meaning in relation to the organism. You will see the
>"facilitator" as doing nothing more than applying disturbances and
>rearranging the environment, without any special effect on the organism.

I agree. The "If... [then]" is correct, I think. And I think it accurately describes the postulates and implications of your ideology. (It does NOT describe the postulates and implications of PCT, in my view.) Now, consider who cares about the truth of this statement. "Nothing more," eh? The "facilitator" (if he/she abides by the rules for success in the endeavor, as I have been attempting to ascertain using principles of PCT -- which basically means that he/she must make a sufficiently correct model of certain reference signals of the other organism) can in fact by "doing nothing more than applying disturbances and rearranging the environment" have the effect of controlling some of his/her perceptions dependent on actions of the other organism, which otherwise would not be controlled. "Without any special effect on the organism" is a big "so what" clause for people who want to control perceptions which depend on the actions of other organisms, and for people who are subjected to disturbances by, and/or whose environments are rearranged by, people who want to control perceptions which depend on their (the facilitatees') actions. I claim that those people include virtually everyone (absolute hermits, if any exist, participate in no social interactions).

You always try to make the point that unless A can make B want what A wants B to want, the interaction is "unimportant" to B. Do you really think that many people care about this sort of "unimportance"? I look around and see that, by and large, people are trying to control perceptions (of their own) which depend on others' ACTIONS, and usually don't care much about others' WANTS. The boss doesn't care what the employee spends his/her paycheck on, as long as the employee performs actions in the boss's conceptual class "doing the job as I think it should be done." You seem to believe that Skinner's conception of "control" of another organism's behavior is an anomaly. I think it is the rule, rather than the exception. I understand that Skinner got the theory wrong -- and I think PCT is right -- but he was investigating a phenomenon of tremendous significance in everyday life: control of one's perceptions which depend on others' actions. He made the error of calling it "controlling" others. It is a common error. One CANNOT PCT-control others, one can PCT- control only (some of) one's own perceptions, and nothing else in the world, including "others." In some recent posts, even you have (verbally) lapsed into the same error, suggesting that "controlling to see another act in way x" = "controlling the other's actions." Regardless, I maintain that human concern with control of perceptions which depend on others' actions is nearly universal; to date, you are the ONLY person I've heard say that it is unimportant to the others on whose actions the controlled perceptions depend.

Sometimes I think you hint that the world would be a much better place if only people stopped trying to get other people to want what they want them to want. Maybe so. A step toward achieving this goal might be to point out that wanting another to want what you want him/her to want generally results in conflict, escalation, and violence (unless the would-be controller gives up), and that there are (sometimes, with limits as prescribed by PCT) ways to get what you want WITHOUT resulting in conflict.

>The reinforcer is no longer seen as having any special effect on the
>organism, other than the effects that any sensory input has.

If an organism acted identically in the presence of reinforcers and nonreinforcers, there would be no way to distinguish them. But there IS a way to distinguish them (at least in principle): the Test for the Controlled Variable. Investigator Skinhead determines that two rats in his lab DON'T have the same Controlled Variable -- call it "satisfying hunger" -- since Alfy gobbles food each day at 3 PM and Bety sniffs at the food and just sits there. Skinhead gets the crazy notion that maybe there is something different about the rats' environments. He checks all the automatic feeding mechanisms for morning and evening chowtimes -- nothing amiss. But then he discovers that a grad student has taken pity on "poor" Bety and has been sneaking in at 2 PM each day and giving Bety as much food as it will eat. So Skinhead publishes a paper (in SCIENCE, of course) allowing as how the environmental history of a rat can alter what counts as a reinforcer in the here-and-now. He gets his Nobel Prize in Physiology/ Medicine. So far, so good. But then Skinhead gets another crazy notion: an organism's environment is IN TOTAL CONTROL of that organism, so "getting the reins" of an organism's environment can allow one to CONTROL the organism in unlimited ways. Unfortunately for Skinhead, PCTers come along and show that environmental history is NOT "in control." The Nobel Prize is not retracted, because the #1 PCTer claims that the environment, living or not, can have "NO IMPORTANT INFLUENCE" on organisms; as crazy as Skinhead's notion of TOTAL environmental control seems, the notion of NO environmental influence seems crazier. Of course, this isn't EXACTLY what the #1 PCTer means, but when he explains what he REALLY means, the Prize committee simply shrugs and says, "OK." About 30 years later, a "whole new approach" (say the instigators) to psychology is adopted far and wide, based (if the truth were known) on the ideas of the #1 PCTer, for which he is given absolutely no credit; some call this "ironic," since, after all, the #1 PCTer still maintains that he could have had "no important influence" on the new psychologists even as he periodically attempts to show that they got their "most important" ideas from him.

>So I can say that I present an organism with a certain environmental
>situation. I cannot say, without aid of a theory, that doing this amounted to
>"facilitation" or "teaching" or any of those abstract terms that carry within
>them a theory of causality, that imply an arrow reaching from the outside
>world into the organism and changing something inside the organism.

I have been attempting for some weeks to couch my claims in, not theory- neutral terms, but PCT terms, with the exception of trying to correlate PCT- derived definitions with folk and/or sociological terms. I have been trying to consider the data NOT in a vacuum, but from the standpoint of PCT -- sans ideology.

>All descriptions that imply an effect inside an organism
>controlled by an external agency are based on the S-R concept of
>behavior and the causal model that supports it.

I do not claim that "an effect inside an organism" can be "controlled by an external agency," except by the employment of overwhelming physical threat (and maybe by threat thereof). I do claim that under certain (PCT-defined) conditions, one person can control his/her perceptions which depend on others' actions. I further claim that explaining the nature and limits of such control is of great significance to many people. I further suspect that explaining the nature and limits of being able to make someone else "arbitrarily" do what you want them to do is of less significance to many people.

>All data concerning such effects, however genuine, can also be interpreted a
>different way under control theory.

I've been taking PCT as a given, and I've been attempting to interpret the data using PCT, certainly NOT using S-R theory... or organismismic ideology... or radical environmentalismic ideology.

A "deceptive dog" story, told to Pat yesterday: Two dogs (same breed and size) in a household were given a treat of leftover meat (one hunk). One of the dogs roared away to the door and began barking like crazy. (The dogs do that when anyone appears at the door.) The second dog followed, barking. Immediately, the first dog ran back to the meat and chowed down. The second dog remained, barking, at the door, long enough for the first dog to get all of the meat.

To netters: Has anyone other than Bill and I really read all the way to here? If you have, do you think it worth reading? Please let me know at your earliest convenience. If I receive no positive replies, I think I'll post no more in this well-trod field, except maybe a "Warning! Live Mines!" sign.

Back to housebuilding!   Greg

Date: Sun Oct 18, 1992 12:24 pm PST
Subject: Debate, Rick's diagrams

[From Dag Forssell (921018-1)]   Greg Williams (921018)

>To netters: Has anyone other than Bill and I really read all the way to
>here? If you have, do you think it worth reading?

Your correspondence has been so voluminous and hard to grasp in detail,
that I cannot say I have "really" read it. I have perused it and have enjoyed it. You have both provided lots of food for thought. I think Greg has provided us all a very valuable service by bringing up the subject of influence and being doggedly persistent. The CSG group is indeed free of intellectual snobbery, and this is a major strength. Another group I can think of has a guru who proclaims the truth from on high and does not allow himself to be questioned or challenged.

It seems to me that Greg is expressing a lot of "real world" concerns in this latest post. PCT practicioners have to deal with these.

Leadership and management is control by definition. But people cannot be controlled from the outside. How do you resolve this obvious management paradox? With threats of violence, of course. But if you are enlightened? You give people information and time to allow them to align their wants with yours. Even Bill Powers as computer expert at the newspaper did not spend all day on PCT, as he might have wanted to. Bill wanted the computers at the paper to function.

The other day, Bill remarked that dogs seem capable of this... As I understood the A vs. B algebra (I did not try very hard), the issue was anticipating the want structure of the other individual. I think a loose interpretation is that dogs too can live by the golden rule. What sets PCTers apart is the recognition of the diamond rule.

Fire! swimming and pain have been brought up as causing reorganization. I fail to see that either has a direct bearing on critical variables (except perhaps holding your breath causing an excess of CO2 in your lungs, which gives a large error signal to the autonomous breathing control system). Here is a chart from my presentation, where I try to distinguish between critical variables and what people talk about.

STRUCTURE OF CHOSEN WANTS

GENETIC --->	INDIVIDUAL --->	SELECTED
REQUIREMENTS	"HUMAN NEEDS"	"WANTS
examples:	examples:	examples:
Food (various nutrients)	Belonging	Perform task
Water	Love	Specific food
Oxygen	Sex	Specific drink
Temperature	"Self-worth"	Rest
Salinity	"Esteem"	Transportation
Calcium (in blood)	Safety, survival	Possessions
Control (within cells)	Control, power	Status
Various hormones	Learning, fun	Relationships
We have some genetic	These requirements	We select wants in

requirements to live and thrive.	we interpret by defining individual "human needs"	order to satisfy our genetic and individual "human needs" as we understand them

On the subject of feelings, my concept is that when we have an error signal, we generate both logic signals, such as for muscle action AND hormone signals, such as adrenalin for "release of energy." Of course it is not quite that simple. Norman Cousins suggested that our brain secretes thousands of hormones.

In my teaching, I suggest that your information content (what you already know) determines (influences?) how you perceive something. Your information content also determines what you choose to want in relation to what you perceive. Your information content finally determines your range of choices of what to do with that error signal. (The three are integrated, of course - that is the point of behavior of perception).

Whether fire!, drowning or torture is perceived, how you perceive it, what you want in relation to it and what you do about the error signal, if any, depends. If you release massive amounts of adrenalin or whatever, that may upset your critical variables and you reorganize. If you don't, you don't.

Greg and Bill: Your discussion on manipulation and influence is indeed worth reading. It has provided some of the meatiest and most worthwhile clarifications of PCT in a long time and has stimulated some carefully developed definitions, which many of us will study and work with for a long time.

Better the mines exploding in earnest debate than questions being ignored.

Swedish poetic rewrite of Icelandic saga: Gunnar has been feuding for a long time with other independent farmers. He is cornered, under attack in his farmhouse which is now on fire. He is in the attic, his wife with him. The string breaks as he pulls it to shoot an arrow. He turns to his wife and asks her to cut some of her long hair and twist him a new string. She refuses, noting that her relatives are attacking. Obviously, they are both about to fry. Says he laconically: "Better listen to the string that broke than never pull on your bow." Idiomatic translation: "Better to have loved and lost than never loved at all."

____

Rick's charts: Is this treatment of the environment an improvement?
____

FIGURE 1: A BASIC CONTROL SYSTEM

```
                              r
                              |
                        ______v_____
                       |            |
                   o-->| comparator |---o
                 p |   |____________|   | e
              _____|_____            _____v_____
             |           |        |           |     Control
             |  sensor   |        | amplifier |     System
=============|___________|========|___________|=================
                   ^                    |           Environment
              _____|_____            _____v_____
             | controlled|        | amplified |
             |  input    |        |  output   |
             | variable  |        | variable  |
             |     i     |<-------|     o     |
             |___________|        |___________|
                   ^
              _____|_____
             |  disturb- |
             |   ance    |
             |     d     |
             |___________|
```

```
      |  = down        ^   = up         o-->  = connection and
      v    arrow       |     arrow      |       side arrow
```

FIGURE 2: A THREE LEVEL HIERARCHY OF CONTROL SYSTEMS

Only the left half of figure 2 shown here. The right half extends the picture to the full four "stacks" of interconnected, interrelated control systems.

```
                    r                                   r
                    |                                   |
Level 3       ______v_____                         _____v______
             |            |                       |            |
         o-->| comparator |---o               o-->| comparator |---o
       p |   |____________|   | e           p |   |____________|   | e
    _____|_____          _____v_____     _____|_____          _____v_____
   |           |        |           |   |           |        |           |
   |  sensor   |        | amplifier |   |  sensor   |        | amplifier |
   |___________|        |___________|   |___________|        |___________|
         ^                    |               ^                    |
         |                    |               |                    |
         o------------------- | --------------o------------------- | --------
         |                    |               |                    |
         |          o---------o-------------- | --------o----------o---------
         |          | r                       |         | r
Level 2  |    ______v_____                    |    _____v______
         |   |            |                   |   |            |
         o-->| comparator |---o               o-->| comparator |---o
       p |   |____________|   | e           p |   |____________|   | e
    _____|_____          _____v_____     _____|_____          _____v_____
   |           |        |           |   |           |        |           |
   |  sensor   |        | amplifier |   |  sensor   |        | amplifier |
   |___________|        |___________|   |___________|        |___________|
         ^                    |               ^                    |
         |                    |               |                    |
         o------------------- | --------------o------------------- | --------
         |                    |               |                    |
         |          o---------o-------------- | --------o----------o---------
         |          | r                       |         | r
Level 1  |    ______v_____                    |    _____v______
         |   |            |                   |   |            |
         o-->| comparator |---o               o-->| comparator |---o
       p |   |____________|   | e           p |   |____________|   | e
    _____|_____          _____v_____     _____|_____          _____v_____
   |           |        |           |   |           |        |           |
   |  sensor   |        | amplifier |   |  sensor   |        | amplifier |
===|___________|========|___________|===|___________|========|___________|===
         ^                    |               ^                    |
    _____|_____          _____v_____     _____|_____          _____v_____
   | controlled|        | amplified |   | controlled|        | amplified |
   |  input    |        |  output   |   |  input    |        |  output   |
   | variable  |        | variable  |   | variable  |        | variable  |
   |     i     |<-------|     o     |   |     i     |<-------|     o     |
   |___________|        |___________|   |___________|        |___________|
         ^                                    ^
    _____|_____                          _____|_____
   |  disturb- |                        |  disturb- |
   |   ance    |                        |   ance    |
   |     d     |                        |     d     |
   |___________|                        |___________|

Legend:              ^                                 |
                     |                            ---- | ---
              o------o--->                             |
              |      |                                 v
           Connections, arrows                     Crossing paths,
              up & sideways                         arrow down
```

Best to all,   Dag

Date: Mon Oct 19, 1992 3:14 am PST
Subject: The Williams-Powers Debates

[from Gary Cziko 921019.0217 GMT]   Greg Williams (921018) queries:

>To netters: Has anyone other than Bill and I really read all the way to here?
>If you have, do you think it worth reading? Please let me know at your
>earliest convenience. If I receive no positive replies, I think I'll post no
>more in this well-trod field, except maybe a "Warning! Live Mines!" sign.

I think I've read all the posts, although perhaps not all of every one. I've found much of it quite worth reading, and when I'm pressed for time or find something that seems repetitious or uninspiring, I skim through or jump to the next post.

I don't see why anyone should feel reluctant to discuss issues relevant to PCT on the network with anyone else on the network. Indeed I would prefer that the private discussions that have been going between some netters on relevant to PCT all be put on the net so that all may enjoy or disregard as each wishes. I don't see discussions on CSGnet as an imposition on anyone, rather as sharing. As Bill has said, we all have delete keys (or a mouse with which to drag things into the trash bin).

So I say carry on, as long as Bill is willing to continue. I only ask that if there is some particularly important breakthrough or "final" agreement that it be indentified as such in the subject header and posted so that those CSGnetters not interested in the blow by blow but only the final score can be kept informed.

>A "deceptive dog" story, told to Pat yesterday: Two dogs (same breed and size)
>in a household were given a treat of leftover meat (one hunk). One of the dogs
>roared away to the door and began barking like crazy. (The dogs do that when
>anyone appears at the door.) The second dog followed, barking. Immediately,
>the first dog ran back to the meat and chowed down. The second dog remained,
>barking, at the door, long enough for the first dog to get all of the meat.

This is a good story indeed, but did it actually happen? I find it hard to believe, but not living with dogs perhaps I underestimate their powers of perceptual control.

If true, this is a very impressive Umweg solution. For the "smart" dog, he somehow perceives that the shortest way to getting all the meat in his stomach is not to proceed directly to the feast but rather take a detour to the front door and bring along his rival. I wonder if the smart dog is also less dominant in some way so that he would likely lose in a physical face off over the meat.--Gary

P.S. Maybe there really WAS someone at the front door, or at least a passing squirrel or rabbit. Let's get out another chunk of meat and try it again, I say, ever the skeptic.

Gary A. Cziko

Date: Mon Oct 19, 1992 9:28 am PST
Subject: Re: Put your model where your mouth is.

[From Dick Robertson 921019.1140]
New experiments! Yay, I'm for that! Also I have long been curious about why it what would be a good experiment to test that out?
Best, Dick R.

Date: Mon Oct 19, 1992 9:52 am PST
Subject: Address; importance; Dag's diagram

[From Bill Powers (921019.0900)] Chris Love (921019) --

My address: 73 Ridge Place, CR 510 Durango, CO 81301

Greg Williams (921019) --

>What makes most sense to me now, in light of your critique, is to
>hypothesize that the error criteria (to start reorganization) for
>some or all innate critical reference signals can be altered by
>learned systems. That would explain the patriot steadfastly
>refusing to tell state secrets under torture, I think; the
>patriot's learned "honor" (or whatever) system increases the error
>criterion for pain sufficiently that reorganization is not
>triggered up to the point of death.

If my model is correct, torture can cause reorganization, but can't cause any PARTICULAR reorganization. Most people who are tortured eventually do or say whatever will make the torture stop. Perhaps all of them do -- how can we ever know? Perhaps the few who were steadfast unto death simply didn't reorganize to produce the required change in their control systems soon enough. They were, depending on how you look at it, the unlucky ones, considering the actual value of most secrets that prisoners of war know. Returned prisoners who have been tortured and have spilled the secrets or confessed have said time and time again that all people will crack under torture; stories to the contrary are macho military myths which only add to the misery of people who have undergone these experiences. You had better check your genuine social data before believing such myths. The fact that an apocryphal story supports your position is not enough reason to believe it.

>But maybe you think "not experiencing pain" is NOT an innate critical
>reference signal?

It is not. It is a learned goal. The reorganizing system, as I conceive it, would respond to damage, but not to sensory pain. This is why people can learn to seek sensory pain. Autonomic sensing complicates this a bit, but autonomic signals are not part of the hierarchy as far as I can tell.

>>If, on the other hand, you think that all changes are internally
>>motivated and accomplished, you will see the very same actions in
>>the outside world as having a different meaning in relation to the
>>organism. You will see the "facilitator" as doing nothing more
>>than applying disturbances and rearranging the environment,
>>without any special effect on the organism.

>I agree. The "If... [then]" is correct, I think. And I think it
>accurately describes the postulates and implications of your
>ideology. (It does NOT describe the postulates and implications of
>PCT, in my view.)

I'm getting pretty tired of this "ideology" crap, buddy. Let's start breaking down some meanings into PCT terms.

Let's do "importance." DEF: Something is important to a person if the person perceives it and has a reference signal to compare it with. We observe such things by finding controlled variables with reference levels. We discover their importance by disturbing them and finding that an action changes in such a way as to prevent the disturbance from affecting the controlled variable (materially).

```
                             ref signal
                                 |
                     perception  |        error
INFLUENCEE           ---------- COMP ----------
                    |                           |
                    |                           |
             INPUT FUNCTION               OUTPUT FUNCTION
                    |                           |
            --> cont. var.<--env. link<-----   action <---
           |       |                                      |
           |   dist var.                                  |
           |                                              |
           |                                              |
            <-- this is important to the person           |
                                                          |
                  this is not important to the person -->-
```

The influencer acts either by altering the disturbing variable or the environmental link. The result in either case is a change in the action. No alteration in the organization of the control system is required unless control is made impossible. The action simply changes as required to maintain the perception near the reference signal. The "action" label could include many lower-level control systems.

As a result of altering the disturbance or the link, the influencer perceives a change in the action. The influencer can therefore learn to control the action in the above diagram to match the influencer's reference level for it (via perceiving, comparing, and acting as usual). The influencee's action as perceived by the influencer is therefore important to the influencer by the same definition.

The only way for the action to become important to the influencee would be for it to be perceived and compared with a reference signal of its own. But this would mean that both the action and its effect on the controlled variable in the diagram would be independently controlled. This would immediately produce conflict, because for the controlled variable in the diagram to be controlled, the action must be determined by the disturbing variable and the reference signal in the diagram. Any other effect on the action would constitute a disturbance, which the control system in the diagram would resist by altering its output. So the influencee cannot have goals both for the controlled variable and for the action that controls it, if conflict is to be avoided.

So the influencing person can control the action of the influencee, but the influencee cannot have any preference for one degree or sign of action over another. Nor can the influencee control the influencee's own action.

The controlled variable is important to the influencee and is controlled by the influencee but not by the influencer. The influencee's action is important to the influencer and is controlled by the influencer but not by the influencee. This all presumes no change of organization and no loss of control by the influencee.

So far this is straight PCT, is it not?

Dag Forssell (921019) --

That's pretty good work with ASCII diagrams! I have used your first one, simplified, in the reply to Greg above.

Best to all,   Bill P.

Date: Mon Oct 19, 1992 10:00 am PST
Subject: Debate, Rick's diagrams

[From Rick Marken (921019.1000)]   Dag Forssell (921018-1)

The diagrams are beautiful. I am thrilled that you put put all that hard work into it. They are beautiful charts.

Greg Williams (921018)--

>To netters: Has anyone other than Bill and I really read all the way to
>here? If you have, do you think it worth reading?

Haven't read every word but I think it's very worthwhile.

Keep it up.   Best regards   Rick

Date: Mon Oct 19, 1992 4:39 pm PST
Subject: oxymorons

[From Wayne Hershberger 921018]

(Bill Powers (921011) re: Wayne Hershberger (921010) --

WH   I would prefer something like this: "It is that the world we experience directly is ALREADY being shaped by the perceptual processes even as we are experiencing it."

WP   What are these processes and how do they do this shaping?

Reflection and refraction of light for one thing, synaptic transduction for another, perceptual input functions, for a third; all of these and more. The perceptual process involves an ecological dipole. Perception is NOT simply a process of transporting a representation of some putative conceptual reality comprising one end of the dipole (the environmental pole) into the other end (the organism pole). Such a conceptualization (representationalism) begs the fundamental question of perception, which is the realization of the perceptual world in the first place. For example, consider two light emitting diodes (LEDs) moving in phase in the dark, one vertically, the other horizontally, like this (i.e., conceive it thus):

```
                  ^
                  |
                  |
                  |
                  v <-------->
```

Doing this, Gunnar Johansson found that one perceives diagonal motions. The two lights are seen as separating diagonally as the PAIR moves along the orthogonal diagonal, from lower left to upper right; and then as the PAIR moves diagonally back to the lower left the two lights are seen to approach each other, diagonally. Describing the LED's motion as vertical and horizontal is a conceptual convenience. And it is realistic. But this conceptual reality doesn't account for the diagonal motions that are perceived. And for the same reason, it doesn't account for the perception of vertical and horizontal motion (e.g., when the room lights are on), either. Why should it?

WH   Thus, conceptual models of perception merely assert an equivalence between perceptual and conceptual realizations. On this point I trust we are all agreed.

WP   I would agree if I knew what you meant. What do you mean by "perceptual" and what do you mean by "conceptual?" Pretend I really don't know; you will be right. Are concepts not perceivable?

The phenomenon described above illustrates the difference. One can conceptually represent (i.e., re-present) a perceptual reality as I did above in describing the LED's motion as vertical and horizontal--that is, I have described the motion (i.e., conceptually realized it) as it looks (i.e., is perceptually realized) in the light. But this equivalence (between the perceptual and conceptual realizations) is an accident of the room lights being on; when the room lights happen to be off, the equivalence vanishes. However, this is not to say that the perceptual realizations (vertical/horizontal motion or diagonal motion) are themselves an accident of the room lights being either on or off. The perceptual realization DEPENDS LAWFULLY UPON the status of the room lights- -same brain, different input. The input, of course, is neither "diagonal motion" nor "vertical/horizontal motion," but rather something from which these alternate realizations may be achieved. Perceptual realizations are not caused, they are achieved.

A conceptual realization may be represented (re-presented) perceptually only by controlling the perceptual world. That is, in order for a conception to be perceived the conception must first comprise reference signals. For example, once an architect has conceived a structure and represented that conception in blueprints, a carpenter can build the structure for all to perceive.

WH   The expression _conceiving a percept_ is an oxymoron because a percept conceived is a concept not a percept.

WP   Does this suggest a hierarchical relationship? That is, perceiving processes produce percepts, from which a conceiving process can produce a concept?

The difference is not just the level in the hierarchy; conceptualizations involve a lot of imagination.

Warm regards, Wayne

Date: Mon Oct 19, 1992 8:40 pm PST
Subject: flush

from isaac n. kurtzer(921910)

i think that above symbol is correct for dating mail? i'll check back at the the room.

----

to w.t. powers & g. williams: about your arguments about control of others- i have enjoyed the discussions, but understand control as tight which eliminates the idea of control of others; unless sloppy and at-a-distance "control" counts as CONTROL (sorry if its awkward). also i think a FUNDAMENTAL issue is that for Dr. Nil-point to "control" a obviously stupid rat (sarcasm) he can'tkill it, build a house, or ride a bicycle; he IS ALSO RESTRAINED; bourbon wrote a paper on this and i feel that is very much ignored by people that pride themselves on being controllers. what i see from them is their wasting their time (possibly dedicating their lives or sacrificing their culture--in fact, that is mandatory) to set up the conditions necessary for the "controlled" to produce actions it doesn't care about anyway. what a waste and delusion.

----

to anyone (esp. bourbon): for my last tech. writing assignment i was required to write an expanded definition on a mechanism/device, so i did mine on the unsung regulatory device- the toilet, yes the toilet. porcelin homeostasis!!! i have to amuse myself somehow; oh yeah, dr. bourbon, i have talked to gaylord and if the undergraduate perception class doesn't open and since i refuse to take either cognitive or animal learning, and my extremely FAT A in his class doesn't hurt, i may be able to take grad. perception as a undergraduate!! see you at michelle's proposal.

i.n.kurtzer

Date: Mon Oct 19, 1992 11:36 pm PST
Subject: things with ray...and dag

Hello Dag,

A lot has happened since we last communicated. My Dad passed a little over a week ago, and I'm going through that overwhelming grief that few of us are spared. My Dad was a fine man, some say simple, but high-school educated, WWII Vet, moved to Arizona in the 50's with his young bride and $40, ended up buying a house and raised 3 boys (two cops and a teacher). He was proud of his boys and, being a "common sense" kind of factory-guy at heart, he was very proud of some the of the work I was doing, even though it was difficult for him to understand.

On the other hand, Dag, although I'm WAY behind on the Net, I read your stuff. I was happy to hear about the Deming connection...from the lack of Net replies on that, I don't think some of the others realize how truly significant that recognition is to the industrial world. I have many thoughts on that I'd like to Net-out, but it's still difficult, both time-wise and emotional-wise.

I hope you didn't find this note to be a "downer"; I'll be in touch soon when I have a better handle on things.

Sincerely, Ray

Ray L. Jackson 602-963-6474
3613 W. Saragosa St.
Chandler, Az 85226
attmail.com!rljackson

Date: Tue Oct 20, 1992 8:27 am PST
Subject: greg and bill stuff

David Goldstein 10/15/92

I have been reading your interchanges with interest. The thought occurs to me that a comparison of an HPCT versus a behavior modification approach for a particular case might be relevant to your discussions. It seems that Greg is saying that behavior modification approaches are common and work pretty well in everyday life and that a behavior modification approach educated by HPCT would be even more effective.

I work at a residential treatment center for children ages 12-17. It is an unlocked setting. Sometimes a resident will leave center grounds ("go AWOL"). Recently, a resident was tragically killed during an AWOL. The problem is how to reduce AWOLS among residents to zero.

A behavior modification approach: (a) Consequences have to be identified which when presented after an AWOL to a resident will result in the frequency of AWOLS decreasing. The consequences we have used are: deduct points (relates to money), reduce status level (relates to degree of supervision and activities allowed), conduct a special meeting which includes the resident, treatment team members and administrators (ETRE meetings) which can result in discharge, a special program, verbal warnings and lectures.

(b) Arrange the environment before an AWOL occurs which will act to prevent AWOLs from happening. The "stimulus control" efforts we have used are: personal physical restrainsts when staff judge a resident to be in a state which poses a risk to self/others, verbal statements to residents to stop or return, verbal reminders to residents reminding them of the consequences.

It is obvious from the fact that residents are still going AWOL that the above measures are not controlling the level of AWOLs to the desired level of zero. The measures we are taking might be acting to reduce the AWOL frequency but it would be hard to prove. I would be interested in hearing what suggestions anyone has based on HPCT which would improve the control over AWOLs.

Thanks in advance for your attention to this matter.

Environmental stimuli have to be identified which when present will

Date: Tue Oct 20, 1992 9:43 am PST
Subject: Significant agreement!

From Greg Williams (921020)   >Gary Cziko 921019.0217 GMT

>I don't see why anyone should feel reluctant to discuss issues relevant to
>PCT on the network with anyone else on the network. Indeed I would prefer
>that the private discussions that have been going between some netters on
>relevant to PCT all be put on the net so that all may enjoy or disregard as
>each wishes. I don't see discussions on CSGnet as an imposition on anyone,
>rather as sharing. As Bill has said, we all have delete keys (or a mouse
>with which to drag things into the trash bin).

But such "sharing" can become a burden on busy people, especially if it appears that not many other netters are very interested in it. So far, I count less than 10 "keep it up" posts. So I'll keep this short. I doubt that I'll continue unless I feel I've been considerably maligned.

GW>>A "deceptive dog" story, told to Pat yesterday...

>This is a good story indeed, but did it actually happen? I find it hard to
>believe, but not living with dogs perhaps I underestimate their powers of
>perceptual control.

It was told as matter-of-fact. But it WAS anecdotal, as they say. Grad students, here's the research project you've been awaiting (just try to get it past the animal research ethics committee!)

>Bill Powers (921019.0900)

>I'm getting pretty tired of this "ideology" crap, buddy. Let's start
>breaking down some meanings into PCT terms.

>Let's do "importance."

Again I repeat that I have no problems with your definition. And again I say that many people attach great importance to the outcomes of social interactions involving persons who are controlling their perceptions which are dependent on others' actions, regardless of whether the interactions are "important" as you define it (that is, involving reorganization of the control structure of someone upon whose actions another's controlled perception depends) or are "unimportant." I'm saying that your definition is simply not relevant for people who study such interactions or for people who are actually involved in such interactions.

>The influencer acts either by altering the disturbing variable or the
>environmental link. The result in either case is a change in the
>action. No alteration in the organization of the control system is
>required unless control is made impossible. The action simply changes
>as required to maintain the perception near the reference signal. The
>"action" label could include many lower-level control systems.

I agree.

>As a result of altering the disturbance or the link, the influencer
>perceives a change in the action. The influencer can therefore learn
>to control the action in the above diagram to match the influencer's
>reference level for it (via perceiving, comparing, and acting as
>usual). The influencee's action as perceived by the influencer is
>therefore important to the influencer by the same definition.

I agree.

>The only way for the action to become important to the influencee
>would be for it to be perceived and compared with a reference signal
>of its own.

I agree, but it can become important RETROSPECTIVELY: "Ooops! I was wrong to do that!" And it can be important TO THIRD-PARTY "OBSERVERS" (such as sociologists and police officers).

>But this would mean that both the action and its effect on
>the controlled variable in the diagram would be independently
>controlled. This would immediately produce conflict, because for the
>controlled variable in the diagram to be controlled, the action must
>be determined by the disturbing variable and the reference signal in
>the diagram. Any other effect on the action would constitute a
>disturbance, which the control system in the diagram would resist by
>altering its output. So the influencee cannot have goals both for the
>controlled variable and for the action that controls it, if conflict
>is to be avoided.

>So the influencing person can control the action of the influencee,
>but the influencee cannot have any preference for one degree or sign
>of action over another.

Non-retrospectively, non-third-party, I agree, although I still think that the phrase "control the action" is unfortunate, since the "influencer" is actually ONLY controlling his/her OWN PERCEPTIONS.

>Nor can the influencee control the influencee's own action.

I agree. People can only control (some of) their own PERCEPTIONS, not actions.

>The controlled variable is important to the influencee and is
>controlled by the influencee but not by the influencer. The
>influencee's action is important to the influencer and is controlled
>by the influencer but not by the influencee. This all presumes no
>change of organization and no loss of control by the influencee.

With your technical definition of "importance," here, I agree, again with the caveat about the influencee's action being "controlled" by the influencer.

>So far this is straight PCT, is it not?

It is straight PCT plus a PCT-based definition of importance.

Gary, this the kind of significant agreement between Bill and I which you asked to be noted. Please excuse me now while I go build some walls.

Your buddy,

Greg

Date: Tue Oct 20, 1992 10:12 am PST
Subject: Statistics, perception

[From Rick Marken (921020.0930)] Bill Powers (921013.0930) --

>It strikes me that one problem with "residuals" and all that is simply
>that the wrong model is used (as you say). Is there anything to
>prevent you from doing statistical manipulations using a closed-loop
>model instead of an open-loop one? In fact, isn't that pretty much
>what we do, although informally? We're trying to fit a linear model to
>the data to obtain the minimum least-squares error of prediction,
>aren't we? The only difference is that our linear model embodies a
>closed loop.

Yes. As I said to Martin, I think it would be great if he could show HOW he would go from observation of the data to modification of the model. My impression (which Martin has yet to dispel) is that Martin was suggesting that you COULD improve the control model by analyzing the data in terms of an open loop model (and Martin's statistical model IS an open loop model). If Martin had something else in mind, then I wish he would clear it up for me. As I said, I think it would be an enormous (and, I think feasible) contribution to PCT if he could develop an analytic technique for going from discrepencies between model and actual behavior to revision of the parameters of the closed loop model to improve the fit.

Wayne Hershberger (921010) --

>Perception is NOT simply a process of transporting a representation of
>some putative conceptual reality comprising one end of the dipole (the
>environmental pole) into the other end (the organism pole). Such a
>conceptualization (representationalism) begs the fundamental question of
>perception, which is the realization of the perceptual world in the first place.

What is your model of perception, Wayne?

Here is my model:

EV --->S--->PF--->PS

where EV is an evironmental variable, S is a sensor, PF is a perceptual function and PS is a perceptual signal. The EV is known only in terms of our models of physics. But whatever EVs REALLY are, they impinge on S (based on our models of physiology) which transforms the EV into neural signals (the physiological model again) that enter a neural netork that acts as a computation device (PF) that converts the input neural signals into an output neural signal, the perceptual signal (PS) -- all this is based on the physiological model. I imagine that it is PS that IS the experienced perception (this is the PCT mind model). In your example of "diagonal movement", PS IS the perception of diagonal movement -- constructed from the sensory inputs that are ultimately caused by the horizontal and vertical EV movements. What else is needed here -- other than the delineation of how PS results from EV -- ie. other than the model of PF and S? What is your model of perception? How does it work?

Gary Cziko --

Thanks for sending the note to Ari. It was very clear and helpful as usual.

By the way, I have not heard a peep from Psycholoquy about my Blindmen paper. Better start thinking about the next place to send it.

Best regards Rick

Date: Tue Oct 20, 1992 10:43 am PST
Subject: greg and bill stuff

[From Rick Marken (921020.0950)] David Goldstein (921015?) --

>It is obvious from the fact that residents are still going AWOL
>that the above measures are not controlling the level of AWOLs to
>the desired level of zero. The measures we are taking might be
>acting to reduce the AWOL frequency but it would be hard to prove.
>I would be interested in hearing what suggestions anyone has based
>on HPCT which would improve the control over AWOLs.

I'm sorry. I can't resist.

The obvious HPCT answer is "kill 'em". The only fool-proof way to control a control system is make it into a cause-effect system -- ie. a corpse.

What, the staff didn't like that suggestion? Hmm. Must be a bunch of liberals.

Best Rick

Date: Tue Oct 20, 1992 10:49 am PST
Subject: Realizations; Control of controller

[From Bill Powers (921020.0900)] Wayne Hershberger (921020) --

WP>> What are these processes and how do they do this shaping?

> Reflection and refraction of light for one thing, synaptic
>transduction for another, perceptual input functions, for a third;
>all of these and more.

How do you know about "light" and its reflection and refraction? It seems to me that "light" can never be made into a perception, although it can be conceptualized (imagined).

>Perception is NOT simply a process of transporting a representation
>of some putative conceptual reality comprising one end of the
>dipole (the environmental pole) into the other end (the organism pole).

I agree. As I conceive it, perception is a process that produces percepts, which we deduce to arise from phenomena such as light etc. which are not represented in perception. All we know consists of percepts or realizations.

>... begs the fundamental question of perception, which is the
>realization of the perceptual world in the first place.

I assume that the perceptual world IS the realization; that is, we do not experience a something and a realization of the something as two separate things. We experience only the realization, not what it is a realization of.

>For example, consider two light emitting diodes (LEDs) moving in
>phase in the dark, one vertically, the other horizontally ...
>one perceives diagonal motions. The two lights are seen as
>separating diagonally as the PAIR moves along the orthogonal
>diagonal, from lower left to upper right; and then as the PAIR
>moves diagonally back to the lower left the two lights are seen to
>approach each other, diagonally.

There are three objects that can be perceived: one light, the other light, and an object made of two lights. If the brain assumes a frame of reference moving with the centroid, either in the light or in the dark, and if the eyes track the centroid, then the diagonal movement of the centroid can be perceived only if the tracking motion of the eyes or attention can be perceived. Given perfect tracking of the centroid, there would appear to be two lights oscillating antisymmetrically about the centroid. I should think that given a persuasive story, a person could perceive either of the situations you describe even with the room lights on.

>Describing the LED's motion as vertical and horizontal is a
>conceptual convenience. And it is realistic. But this conceptual
>reality doesn't account for the diagonal motions that are perceived.

When seeing the lights in the dark, it may be that the "vertical" and "horizontal" percepts are missing (perfect tracking). If a background can be imagined, stationary with respect to kinesthetic information, those percepts might be restored. If the subject never experiences the lights in a lit room, I doubt whether the vertical or horizontal percepts would ever be experienced. All perception of position or motion is relative, isn't it?

I don't see a need to distinguish percepts from concepts in this case. Doesn't imagination explain what's going on?

>Describing the LED's motion as vertical and horizontal is a
>conceptual convenience. And it is realistic.

Well, it's realistic as a way of describing the movements of a light relative to a background, if you assume that the background is stationary in some absolute framework. The background itself is moving due to the Earth's rotation and orbital motion, and the Sun's motion, etc... When you remove the background it's impossible to perceive or even define absolute position or motion. When you restore the background, you can describe position and motion only in relation to the background, unless you can perceive that the background, too, is moving relative to something else (looking out the porthole of a ship). The background is also a perception.

As far as I can tell, you're talking about percepts all the way.

>One can conceptually represent (i.e., re-present) a perceptual reality as
>I did above in describing the LED's motion as vertical and horizontal--that
>is, I have described the motion (i.e., conceptually realized it) as it looks
>(i.e., is perceptually realized) in the light. But this equivalence (between
>the perceptual and conceptual realizations) is an accident of the room lights
>being on; when the room lights happen to be off, the equivalence vanishes.

Yes, these are two different functions of (presumably) different input conditions. The difference in the functions is in whether the visual frame of reference is anchored to a background or to the centroid of the moving lights. I can think of a generalization of this experiment in which the "background" consists of a third LED (or also a fourth and a fifth ...) in a darkened room. Now there is a choice about how to perceive the situation. The third LED might be perceived as oscillating on a diagonal line while the other two approach and recede from each other on an orthogonal diagonal line. The ambiguity would be greatest when the motions were small and slow, so information from eye movements and kinesthesia would be minimized. The third LED could in fact also be moving. One advantage of doing the experiment this way would be that imagination would not be involved. Now the difference would be entirely between two different perceptual interpretations of the same scene. If "conceptualization" is simply imagination, it can be removed from the situation. "Imagination" is more clearly definable. If, however, it is a higher-level perception, then a hierarchical model is required to account for the existence of mutually-contradictory percepts, and choice of one rather than another.

>The perceptual realization DEPENDS LAWFULLY UPON the status of the room
>lights--same brain, different input.

The "status of the room lights" is another percept, isn't it? You're saying that the perceptual realization depends lawfully on other perceptual realizations. The lawfulness is perceived in the relative behavior of percepts, with "lawfulness" or "relationship" itself being a higher level of percept, or realization. In my version of the experiment, where no imagination is needed, the switch from one realization to a mutually-exclusive one depends on choice of the higher-level realization -- which "law" is supposed to exist. The appearance of the scene will depend, ultimately, on what stationary frame of reference the brain imagines. The difference [between perception and conception] is not just the level in the hierarchy; conceptualizations involve a lot of imagination.

It looks as if what you mean by "conceptualization" can be modeled in HPCT as a combination of higher-level perceptions and imagination.

Isacc Kurtzer (921019) --

The convention for dates is YYMMDD (year-month-day). This is the European style, which has the advantage that subtracting one date from another always gives a positive result if the subtracted date is earlier than the first. This makes sorting by date easy. I can't think of any other advantage, but we've standardized on it.

The convention I use for dating my communication is as follows:

The writer of the post (me) is identified in the first line in square brackets, with "From" to emphasize who the post is coming from. The date/time of writing is in parentheses, with the time (when I start writing) separated by a decimal point. So today's post from me says [From Bill Powers (921020.0900)].

Then, as I address remarks to or about somebody else's post, I start a new section with that person's name and the date (and time if known) in parentheses, and a couple of hyphens after it in order to complete the aesthetic appearance of advanced typography from an orderly mind, a complete sham.

Behind all this was some vague idea that if everyone used exactly the same notations, it would be possible to write a simple program to strip off headers and search for references to stuff. Of course I'm almost the only one who uses this notation, and anyway I probably wouldn't have actually written or used the program because there would always be somebody who deviated from the exact form (including me when I forgot), which would mess up the whole idea.

> ... about your arguments about control of others- i have enjoyed
>the discussions, but understand control as tight which eliminates the
>idea of control of others; unless sloppy and at-a-distance "control"
>counts as CONTROL (sorry if its awkward).

You can control the actions of a tight control system by applying suitable disturbances to the controlled variable. The disturber in the rubber-band experiment can place the other person's finger over any selected spot, as long as the disturbee maintains the knot exactly over its intended position. This can be done even if the intended position is changed by the disturbee. Have you tried this?

Good news about being able to start graduate courses and skip the B. S. (which does not stand for Bachelor of Science).

Best to all,   Bill P.

Date: Tue Oct 20, 1992 2:52 pm PST
Subject: AWOL; importance

[From Bill Powers (921020.1100)]   David Goldstein (921020) --

There seems to be a conflict at your school regarding keeping residents from going AWOL. If a zero level of going AWOL were the only goal, then a simple way to achieve it (other than Rick Marken's) would be to lock the doors and bar the windows and have armed guards to prevent escapes. This would achieve the desired rate of AWOL incidents: zero.

Obviously, this is not acceptable in relation to other goals: financial, humanitarian, and conceptual (this is not REALLY supposed to be a prison). This unresolved conflict leads to using methods that are ineffective, which result in perhaps a slight reduction in AWOL incidents and leaves other goals not fully satisfied.

There is another problem, as I understand it: your institution operates within the constraints of a coercive legal system. The residents are assigned to you against their will, and are not free to leave whenever they want to. If they escape they will either be returned to your institution or sent to a different one. There are, no doubt, sufficient reasons for doing this in the sense that many people are convinced that it's necessary. We can't change that.

So it seems to me that your problem as people working within the institution is to find the least conflicted way of operating within the external constraints imposed by the legal system. The most important conflicts are not those between staff and residents, but those between and within staff members. To eliminate those conflicts, you must first find out what they are: what goals are behind each method that is used, each rule, each principle? Then you must ask how effective the actions taken to achieve the goals are, and whether each action tends to satisfy ALL goals or tends to satisfy SOME goals while exerting effects away from other goals.

Consider the goals involved in the use of operant methods. If operant methods work, why not use them in the most effective way? The most effective way is to make sure that residents can gain access to neither food nor water unless they behave as you want them to behave. No matter how hungry or thirsty they get, you must not reinforce them until they perform continually closer to the standards you have set. You must of course make sure that they can't give each other any food or water, or find a way to get them themselves, or escape entirely from the conditioning situation. You can put any conditions you like on getting reinforcements: the residents must speak politely to each other and the staff, use good middle-class English, not use words or phrases offensive to any religion, sexual custom, or place of origin, study a certain number of hours per day, sit quietly when not usefully engaged, and keep their quarters neat and clean. The alternative is to die of starvation or thirst. This method will work as long as the residents are in the institution.

If you reject using an effective method, then you must have reasons for not doing the things that are required. These reasons contain or are related to goals that are in conflict with the goal of achieving the above list of ideal behaviors (or whatever your list is). By examining each thing that is done in this way, simply asking what would be require to make it work with maximum effectiveness, you can uncover all the conflicting goals that prevent using any approach in its most effective form.

When you have truly laid out the list of conflicting goals, you will realize the true magnitude of the problem, and will begin to reorganize. You may find that you must prioritize your goals -- it may be more important to prevent physical violence, for example, than to get residents to avoid offending the staff with their language. You may decide to change your definitions of desirable behavior. You may eliminate the offense called AWOL by making leaves freely available (to those who are not absolutely and inescapably confined). You may decide to work out a new system entirely, in conference with the residents, that minimizes all kinds of conflicts. There are many creative solutions that will occur to you once you have convinced yourself that your whole approach to the residents is in a state of severe internal conflict.

Greg Williams (921020) --

\>>The only way for the action to become important to the influencee
\>>would be for it to be perceived and compared with a reference
\>>signal of its own.

\>I agree, but it can become important RETROSPECTIVELY: "Ooops! I was
\>wrong to do that!"

The only way this can happen is for the action to disturb some other perception in an uncorrectable way. The action itself is not perceived, normally, and can never have a reference level. The effect on the other perception is the first thing you notice; it may take some time to track down the action responsible, because normally we pay no attention to our actions. This is simply a case of conflict, in which the action that controls one perception makes control of another impossible. The fact that this takes place over time means only that the solution may have to be found in a higher level of control. But only the perceptions are important to the person; the actions causing the problem will be changed without hestitation, as required. The actions are not important. Only perceptions are.

\>And it can be important TO THIRD-PARTY "OBSERVERS" (such as
\>sociologists and police officers).

So, they have their own perceptions, and try to control for them. So what? This still doesn't make actions important to the person using them for controlling perceptions.

\>And again I say that many people attach great importance to the outcomes

>of social interactions involving persons who are controlling their perceptions
>which are dependent on others' actions, regardless of whether the interactions
>are "important" as you define it (that is, involving reorganization of the >control structure of someone upon whose actions another's controlled
>perception depends) or are "unimportant." I'm saying that your definition is
>simply not relevant for people who study such interactions or for people
>who are actually involved in such interactions.

My definition is precisely relevant to just such people. Their perceptions are important to them, if there are reference levels for them, but their actions are not. The person who studies interactions considers interactions important to perceive. The actions through which such perceptions are achieved -- going to the library, interviewing people, taking pictures, taking notes -- are unimportant and will be changed as circumstances require. People actually involved in interactions consider their own perceptions (of themselves, the other, and the relationships between them) important, but will take whatever actions are needed to maintain the perceived interaction in the state they want. There's no preference for any particular action; all that matters is the state of the perception.

Please note:

> ... "important" as you define it (that is, involving reorganization of the
>control structure of someone upon whose actions another's controlled
>perception depends) ...

At least try to argue against me instead of a straw man. I gave this definition yesterday:

>>"DEF: Something is important to a person if the person perceives it and
>>has a reference signal to compare it with. "

If you like, I will expand it slightly: " ... and tries to control it." Although that is not always the case -- one may not know how to try to control it. The definition above covers it, I think.

Best,   Bill P.

Date: Tue Oct 20, 1992 3:06 pm PST
Subject: need help!

from Ed Ford (921020:1315)

Gary, I tried to get a list of those on the net via your instructions and it didn't work. Would you send the instructions again. Thanks, Ed

Ed Ford

Date: Tue Oct 20, 1992 6:17 pm PST
Subject: AWOLs

[FROM: Dennis Delprato (921020)]

>David Goldstein

>I work at a residential treatment center for children ages 12-17.
>It is an unlocked setting. Sometimes a resident will leave center
>grounds ("go AWOL"). Recently, a resident was tragically killed
>during an AWOL. The problem is how to reduce AWOLS among residents
>to zero.

>A behavior modification approach: (a) Consequences have to be
>identified which when presented after an AWOL to a resident will
>result in the frequency of AWOLS decreasing. The consequences we
>have used are: deduct points (relates to money), reduce status
>level (relates to degree of supervision and activities allowed),
>conduct a special meeting which includes the resident, treatment

>team members and administrators (ETRE meetings) which can result
>in discharge, a special program, verbal warnings and lectures.

>(b) Arrange the environment before an AWOL occurs which will act to
>prevent AWOLs from happening. The "stimulus control" efforts we
>have used are: personal physical restraints when staff judge a
>resident to be in a state which poses a risk to self/others, verbal
>statements to residents to stop or return, verbal reminders to
>residents reminding them of the consequences.

>It is obvious from the fact that residents are still going AWOL
>that the above measures are not controlling the level of AWOLs to
>the desired level of zero. The measures we are taking might be
>acting to reduce the AWOL frequency but it would be hard to prove.

Brief comments from my viewpoint re. modern behavioral theory and practice:

Solutions based on (a) are generally very ineffective. Those based on (b) as described are but dimly justified by behavioral theory; they are more folk solutions. To label them "stimulus control" interventions is rather gratuitous.

The above do not recognize more multivariate developments in behavioral therapy (see Goldiamond, Behaviorism, 1974, 2, 1-84; simplification by Delprato, J. Behav. Ther. & Exp. Psychiat., 1981, 12, 49-55; example of elaboration by Delprato & McGlynn, ibid., 1988, 19, 199-205). I have pretty much quit following this literature but get the impression that it is still being developed.

If I were working on this case, I would first examine closely the complainants, realizing that *all* clinical cases require at least one complainant. Can the complaining be changed? YES. Send the identified client away. NO. There are other systemic considerations, but assume we *do*, for the present, have to deal with AWOLs.

1. What can we help the client to do that will increase his acceptability to others? Leaving the grounds can be used as a reinforcer--with proper planning, gradations, etc.--for performance of socially acceptable behaviors, including academic ones.

2. Why does a client want to leave? Sounds like (a) facility is less reinforcing than outside and/or (b) facility is more aversive than outside. What can be done to make ordinary activities more positively reinforcing and less aversive? Why should the client want to stay there? Of course, it is a tough problem because so many factors operate outside and in settings such as these to keep routine interactions anything but "rewarding."

A constructional approach (Goldiamond's term) suggests we concentrate on building new behavioral repertoires (ah that lingo) rather than on eliminating behaviors. Furthermore, we are advised to begin with what the client gives us; there must be something there from which we can help the development process. What does the client like to do? Read comic books? Fine, this can be incorporated into a plan that point to development of personally and socially acceptable behavior patterns.

Note that there is virtually no hint of constructional interactions in (a) and (b) above. Both are all eliminative. The basic problem with eliminative efforts (show person what *not* to do) is that they do not specify what *to do.* I know not to do a, b, c, d, e, ad infinitum. So what do you want me to do? Why are you keeping it a secret? Furthermore being "still, quite, and docile" is not what I need to know. If a dead person can do it, it is not useful for actually living people. So constructional interventions do require a bit more effort on the part of those in charge; they have to come up with *alternatives* that they find suitable. But, hey, this is fair. They are the representatives of the complainants (even when the complainant is the client).

In their simplest form, constructional therapies are based on the infamous positive reinforcement, but where the reinforcers are more ecological rather than contrived by way of the equally infamous "deprivation schedules." The idea is that the world works best (witness USSR) when there are systematic consequences to response occurrances as opposed to "freely avaliable reinforcers." (As in few of us go through life with incomes that are independent of our efforts--no workee, no payee.) Actually, control system theorists may find Goldiamond's (1974) paper of interest. He wrote it against the background of ethical and legal concerns that came up in the 1970s over "behavior control." Goldiamond suggested the best way to handle such matters was to begin (in U.S.A.) with the Constitution. Not a bad idea.

Despite my distinct lack of enthusiasm for the postulates of behavioral therapy, I do not find superior technologies at present. I do find that knowledge of contemporary behavior therapy is not widespread. The solutions implied by (a) and (b) above hark back to the 1960s. I suppose this might show, in part, that behavior therapists do not do a very good job of communicating new developments in their field--apart from the incorporation of (largely outmoded information processing) cognitive theory.

Dennis Delprato
Dept. of Psychology
Eastern Mich. Univ.
Ypsilanti, MI 48197
Psy_Delprato@emunix.emich.edu

Date: Tue Oct 20, 1992 9:22 pm PST
Subject: control?

to w.t.powers: yes, i have seen and participated in the rubber band demo., what it shows to me is that people control perceptions by varying their actions (the actions are accidental and obligatory!). o.k. but "control" of others i still maintain is poor to futile unless i'm willing to sacrifice a great deal (which isn't unheard of). do you disagree? not to argue about semantics, but does this "control" of others meet the criteria (what criteria? the bill et.al. criteria) for control. maybe it is just good to sloppy influencing (i think their is a difference, actually i think it was mark who wrote the paper on that)

oh yeah, [from isaac kurtzer (921020.2350)

this is really minor i realize, but this control-of-others-jack necessitates loss of "control" by the controller (i.e. self-sacrifice).

who watches the watchman?

a steelworker, a bricklayer, saved! i.n.kurtzer

Date: Wed Oct 21, 1992 5:11 am PST
Subject: Strawbuddy?

From Greg Williams (921021) >Bill Powers (921020.1100)

GW>>And it can be important TO THIRD-PARTY "OBSERVERS" (such as
GW>>sociologists and police officers).

>So, they have their own perceptions, and try to control for them. So
>what? This still doesn't make actions important to the person using
>them for controlling perceptions.

I agree. I'm just (well, it is a pretty big "just"!) claiming that the person using his/her actions to control his/her perceptions (the "influencee") can retrospectively consider having used those actions as being important. And that the "influencer" considers those actions important, and not just retrospectively. And that third-party students of social interactions consider those actions important. (And that, come to think of it, DURING an interaction, a person might consider his/her actions "important" in a different sense from Bill's "importance" -- that is, requiring concern/attention (perhaps for later mulling-over), but not control (i.e., not explicitly setting a reference levels for particular hand movements, but, in some sense, noting movement patterns he/she considers "important").

>There's no preference for any particular action; all that matters is the state
>of the perception.

I agree that often (though not necessarily always; see above), "all that matters" (meaning all that is controlled) is "the state of the perception." But many people will simply say, "Big deal." Those people will include all who, unlike Bill, attach importance to interactions wherein some parties control for some of their perceptions which depend on the actions of other parties.

Ideology comes in if Bill tries to go from an "is" (his definition of "importance") to an "ought" (claiming that "importance" as he defines it should be ALL-important to participants in such interactions, retrospectively or not, and to third-party investigators of such interactions). If he tries to do this, I claim that he will "turn off" a lot of folks to PCT ideas -- folks like the 40-year-old accountant who is glad that he learned the multiplication table in school, and like the little old lady who realizes a year after her roof was "repaired" that it wasn't.

>Please note:

GW>> ... "important" as you define it (that is, involving
GW>>reorganization of the control structure of someone upon whose
GW>>actions another's controlled perception depends) ...

>At least try to argue against me instead of a straw man. I gave this
>definition yesterday:

>"DEF: Something is important to a person if the person perceives it
>and has a reference signal to compare it with. "

>If you like, I will expand it slightly: " ... and tries to control
>it." Although that is not always the case -- one may not know how to
>try to control it. The definition above covers it, I think.

I brought up reorganization because you had, several posts ago, said (in line with your recent definition of "importance") that B disturbing A or altering A's environment is "unimportant" to A if A maintains control throughout, and that B disturbing A or altering A's environment would be "important" to A only if A lost control, becoming conflicted, and therefore reorganizing.

To be explicit, I claim that many people consider as highly important social interactions of the type wherein B controls for some of his/her perceptions which depend on some actions of A, REGARDLESS of whether A's actions during the interaction are important (your definition) to A or not, and REGARDLESS of which of the four types of control in my summary of several weeks back is being used by B, meaning REGARDLESS of whether A reorganizes or does not. If you claim that those people SHOULD NOT consider as important any such social interactions wherein A's actions are not important (your definition) to A during the interaction (so A does not reorganize), then PCT is going to be a very tough sell with respect to what those people consider to be important to them.

Bill, who am I to tell you that you should perceive a problem with PCT being a tough sell? That I believe there IS a problem with it is MY ideology. There is no perceived problem for you if you believe that you are right about what everyone SHOULD believe to be important, and that virtually everyone else is wrong, even if you believe PCT ideas deserve wider attention... until you begin to perceive people shrugging and walking away when you tell them what they should and shouldn't believe to be important. And then I won't have to tell you that you should perceive a problem; you'll be perceiving it.

Best,   Greg

Date: Wed Oct 21, 1992 10:27 am PST
Subject: AWOL's at San Pablo

from Ed Ford (921021:0945)

Jim Graves is a good friend of mine, a former student of mine and a present member of our PCT monthly discussion group. Everyone in the group works in tough social environments, such as sexual abuse homes, mental facilities, corrections and probations, acting out kids in schools, the homeless, groups, etc. We discuss how PCT can be applied in the various situations in which the members find themselves. Jim is head of staff training and development and supervisor of counseling at San Pablo, a residential treatment center for male juveniles 12 through 17. They get most their referrals from the state and courts. I had him read Bill's answer to David (somehow I missed your original message David) and he had some interesting remarks.

First, my own thoughts. What we try to do in our PCT group is THINK as control theorists when looking at how to better apply PCT to our individual situations. You'll notice in his remarks that he perceives both his staff and the juvenile residents as living control systems and he creates and continually changes his methods based on the effectiveness of his ideas. Another benefit of thinking as a control theorist is that it frees you up from much of the behvioral modification and manipulation that goes on in the residential treatment culture. The important issue here is that Jim not only considers the various hierarchy of reference signals (values and beliefs, priorities, standards, criteria and choices) when interviewing and working with the residents, but he also considers their perceptual systems. Obviously he can't control their perceptions. However, he tries to create an environment that would most appeal (influence, if you will) to the resident's reference signals which in turn would directly relate to how strong the juvenile's signals are to improve his life. I think this is part of the genius of what Jim (and others at the center) have done. The more appealing the environment (read loving, accepting, but not weakness), the more the juvenile will be working against what he wants, namely love, acceptance, consistent and fairly enforced standards (read structured environment), respect, something he never experienced at home. Most of all, he will be with a staff that is not only caring, but a staff who actually believes he can make it, again something he has never experienced. Jim's thoughts are as follows:

AWOL's usually occur with new residents who have not agreed with or established standards compatible with the center other than not wanting to be in placement and having the perceptions that they don't have any problems to work on.

It is very important for staff to have common standards or expectations of the residents (read have their own perceptions and reference signals aligned) so that the residents cannot create a split within and among the staff. Both staff and resident expectations need to be realistic and achievable. If expectations are too high, then the residents are discouraged and give up.

San Pablo has reduced AWOL's by beginning to establish a relationship with the prospective resident during the pre-admission interview by using basic Control Theory counseling steps that Ed Ford teaches. The first initial interview is to find out if there is within the prospective resident a willingness (read reference signal) to get a commitment to at least try treatment. San Pablo rarely admits a youth without a commitment. It just doesn't make sense to admit someone who is going to be disruptive to the program. You can't force a commitment. Some who refuse to make a commitment ultimately return because the alternative of returning to a juvenile detention center or a crazy family is worse. But a commitment is critical. We have reduced our AWOL's by 80% (of those who did make commitments) just by involving as many of the staff and successful residents during the pre- admission interview. This establishes relationships with people the youth can relate to once he is admitted to the program. This kind environment gives the youth the needed perception that a commitment will be in his best interest.

During the interview, the youth is given a tour of the facility and the program is explained by a resident juvenile who is doing well in the program. Following admission the staff and resident who interviewed and gave the tour to the new admission, make every effort to make the new resident welcome. Thus, at this very critical time, there is a concerted effort on the part of every one, staff and successful residents, to continue to build on the new relationships begun prior to the admission.

Thus it is important for staff to have their systems concepts, principles, and programs, including their priorities, in line with each other and with the agency. However, without an initial access to the youth's value system, and a commitment to try, and a perception of the staff by the prospective resident as a caring group of people and a perception of some residents as "making" it, there is a high risk of AWOL. Somehow, the belief on the part of the new arrival has to develop that "I can make it here" and it has to develop early on.

There are two problems inherent in residential treatment in today's system. First, kids often come to San Pablo medicated. When kids are taken off medication, many placement agencies perceive that it is time to discharge the youth because there is no longer any "medical necessity." Or, when a youth, full of drugs, stabilizes, many see this as meaning it is time to discharge the youth. Either way, the youth's control system, especially his systems concepts, principles, and program levels, HAVE NOT BEEN ACCESSED in any way. All that has been achieved is drugging up the youth. Second, many placing agencies will discharge kids from treatment just when improvement is beginning. This is a shame because at this point most kids are in a state of reorganization dealing with their higher levels and are beginning to sense the possibility of success and the staff as a vehicle to that success.

Residential treatment requires a long period of time (12 - 24 months) in order for the emotionally and behaviorally disturbed youth to truly reorganize their systems with any degree of success. When discharged too early, they revert to earlier ways of resolving conflict since they haven't built the self-confidence in the successful ways they are trying to build. The success of any program begins with building trusting relationships. Without that, nothing else will work. Then, with a commitment to the program, begins the long, slow process toward evaluating and committing to resolving the many conflicts within the youth's values and beliefs, priorities, standards, and various options or choices open to him. For the youth to have any success, he must ultimate believe that the staff has confidence in his ability to succeed. A respect has to develop for each other as a living control system that has it's own internal goals and desires. Before discharge, the resident has to build the strong and confident belief that he can resolve his problems and live a happy and productive life.

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

Date: Wed Oct 21, 1992 10:52 am PST
Subject: AWOLS, PCT Popularity

[From Rick Marken (921021.1000)] Dennis Delprato (921020) --

>Despite my distinct lack of enthusiasm for the postulates of
>behavioral therapy, I do not find superior technologies at present.

The "superior technologies" of behavior control already exist (as Bill and I noted); they involve the use of overwhelming force; lock-ups with armed guards, complete restriction of access to substances required for life (food, water) and, failing that, a shot through the temple.

You can try all you want to control a control system but, as Powers points out eloquently and clearly in the Powers/Williams debate, your "control" is at best ephemeral (at least, when you are dealing with a control system that is organized the same as the control system that is trying to do the "behavior control"-- ie. a control system that controls the same perceptual world). When a control system tries to control other control systems the typical result is conflict -- unless you just want to see the control system produce an action that is irrelevant to the control system itself (the dog happily puts it's paw in the air to get all that dumb love that it really cares about).

>I do find that knowledge of contemporary behavior therapy is not widespread.

Not nearly as un-widespread as it should be.

>The solutions implied by (a) and (b) above hark back to the 1960s.
>I suppose this might show, in part, that behavior therapists do not do
>a very good job of communicating new developments in their field

Dennis, you seem to believe that there is a good approach to behavior therapy. How could this be? PCT shows that behavior therapy could only make sense if it were an effort to help a person control their own perceptions relative to their own goals. "Behavior" as something seen by the therapist is irrelavant to the therapee -- but, like going AWOL, it may be quite important to the therapist. So maybe the term "behavior therapy" is just misleading. Perhaps it should be called "personal control" therpy -- unless, of course, the goal of the therapy is really to make the therapist feel better.

Greg Williams says:

>Bill, who am I to tell you that you should perceive a problem with PCT being a
>tough sell? That I believe there IS a problem with it is MY ideology. There is
>no perceived problem for you if you believe that you are right about what

>everyone SHOULD believe to be important, and that virtually everyone else is
>wrong, even if you believe PCT ideas deserve wider attention... until you
>begin to perceive people shrugging and walking away when you tell them what
>they should and shouldn't believe to be important. And then I won't have to
>tell you that you should perceive a problem; you'll be perceiving it.

I think this gets at the heart of Greg's complaint; he would really like to see PCT ideas get wider attention. I think he sees a lack of interest on the part of some PCTers (like myself) in finding common ground with psychologists, roboticists, biologists, AIers, ALifers, etc etc -- ie. with others in the community of life scientists who might profit from an examination of PCT ideas.

Those of us who do not seem to "compromise", "see commonalities", etc believe that we are just presenting the PCT model -- we don't feel that there is an agenda to alienate potential friends; but apparently it seems like this is true to some people (like Greg, I think).

I think this is worth a discussion. I take the view that the way to promulgate PCT is to present the model and the research honestly, doing what we can to relate this to existing relevant concerns, but not shying away from explaining the true implications of the model. I believe that attempts to "find common ground" produce the Carver/Scheier approach to PCT -- which ends up using the terminology of PCT but misses the basic point (and succumbs to the causal view of behavior in the end). I realize that the approach I advocate is not a good way to drum up a PCT following -- most people do shrug and walk away simply because they don't see the problem that PCT solves -- or don't get PCT even if they do see the problem. But I don't think it is worth it to compromise the model to try to get recruits -- PCT is neither a religion nor a political party. And every so often someone DOES get stoked on it (like I did). I think we are gaining PCTers who really get it (ie. model- ers) at the rate of about 1 a year now. That's plenty for me. I just don't think you can MAKE people be interested in PCT. When the light does go on in a person, PCT sells itself. It's not worth distorting the model to try to get people interested -- you just end up with people who are really interested in the version of PCT that you made up for their sake. But, I'd like to hear your point of view Greg.

Best regards   Rick

Date: Wed Oct 21, 1992 11:03 am PST
Subject: Importance & problems

[From Bill Powers (921021.0915)]   Greg Williams (921021) --

>>So, they have their own perceptions, and try to control for them.
>>So what? This still doesn't make actions important to the person
>>using them for controlling perceptions.

>I agree. I'm just (well, it is a pretty big "just"!) claiming that the
>person using his/her actions to control his/her perceptions (the "influencee")
>can retrospectively consider having used those actions as being important.

I think you're still missing my point. The only way in which anyone can even know what actions he or she is producing is to perceive them. To "retrospectively consider having used those actions as being important" can mean, under PCT, only that the person experienced a perception dependent on the outputs at the time they were performed, remembered it, and later considered it to be an important perception (i.e., adopted a reference level for it).

>I agree that often (though not necessarily always; see above), "all that
>matters" (meaning all that is controlled) is "the state of the perception."

If you can think of anything else beside a perception that can be controlled (or be important to a person) then you are proposing a different model from PCT.

>Ideology comes in if Bill tries to go from an "is" (his definition
>of "importance") to an "ought" (claiming that "importance" as he
>defines it should be ALL-important to participants in such
>interactions, retrospectively or not, and to third-party
>investigators of such interactions).

According to PCT, nothing CAN be of importance to a person but perceptions. This applies now, later, and to third-party observers. All knowledge of the world comes into the brain in the form of perceptions. There is nothing else to control. The outputs of a person are known to that person only to the extent that they affect that person's perceptions, and only the perceptions affected by the outputs can be controlled. The outputs of a person affect other people's perceptions differently, in general, from the way they affect the person's own perceptions. The action I perceive myself performing is not, in general, the action that others see me performing. And neither my perception nor that of others is a direct apprehension of my outputs or their actual physical effects.

>If he tries to do this, I claim that he will "turn off" a lot of folks to
>PCT ideas -- folks like the 40-year-old accountant who is glad that he
>learned the multiplication table in school, and like the little old lady
>who realizes a year after her roof was "repaired" that it wasn't.

If you would try analyzing these situations in PCT terms, under which all the accountant or the little old lady can do is to control perceptions (including a perception of "gladness" and a perception that is "realized"), perhaps you would see that there is no contradiction. On the other hand, if you simply take appearances at face value and give them their traditional informal non-PCT naive realist interpretation, you will continue to miss my point.

>I brought up reorganization because you had, several posts ago,
>said (in line with your recent definition of "importance") that B
>disturbing A or altering A's environment is "unimportant" to A if A
>maintains control throughout, and that B disturbing A or altering
>A's environment would be "important" to A only if A lost control,
>becoming conflicted, and therefore reorganizing.

My fault. I should have said that I was offering a superseding definition. "Importance" is an ambiguous term. You can say that eating is important to someone, and to prove it cite evidence showing that the person IS controlling successfully for eating. Or you can say that that person wants to eat, but is not able to do so, and that eating is then important because of NOT being controlled. "Important" turns out to be a pretty vague term when you break it out into PCT. By offering my definition I was trying to settle on the first meaning.

It's much simpler to say that people have reference levels for perceptions, and normally maintain the perceptions near their reference levels, and when they can't they reorganize. Words like "importance" are typical of the way we speak of experience and behavior in ordinary language; they point to whatever meanings we have in mind, the meanings shifting with context.

This is an example of the ambiguity:

>To be explicit, I claim that many people consider as highly important social
>interactions of the type wherein B controls for some of his/her perceptions
>which depend on some actions of A, REGARDLESS of whether A's actions during
>the interaction are important (your definition) to A or not ...

Do they view those social interactions as important because they ARE being successfully controlled by the viewer, or because they are NOT being successfully controlled? If the viewer sees exactly the social interaction that the viewer wants to see, or if the viewer is able to act in some non-demanding way to make the social interaction return immediately to the desired state, then it is clear that the interaction is important to the viewer, but also that it does not constitute a problem for the viewer. It seems to me that the social interaction that a viewer would consider "highly important" would be one in which small errors will lead to energetic corrective action. But those same social interactions would be "highly important" in quite a different sense if those corrective actions FAILED. Then they would be important in the sense of threatening the integrity of the system; they would call for reorganization.

Perhaps we can use the term "problem" or "difficulty" or some synonym to refer to perceptions that are important because attempts to control them do not work, and reserve the less specific term "important" to mean simply that the person perceives something, has a reference level for it, and acts to correct any difference (or would do so if possible). Thus to say that something is important to a person tells us that there is a reference level for a perception, but does not tell us whether the person is succeeding at controlling the perception. To say that a person has a difficulty with an important perception implies that attempts to control it are not working, and implies that reorganization is likely to be occurring.

And I think we should avoid further confusing the meanings of words by referring to THE importance of a social interaction or anything else, as if there were some objective standard of importance that is independent of anyone's perceptions or desires.

>There is no perceived problem for you if you believe that you are
>right about what everyone SHOULD believe to be important, and that
>virtually everyone else is wrong, even if you believe PCT ideas
>deserve wider attention... until you begin to perceive people
>shrugging and walking away when you tell them what they should and
>shouldn't believe to be important.

I am not telling people what SHOULD be important to them. I am telling them that what IS important to them is their own perceptions. I am telling them that their perceptions are important because of what they desire those perceptions to be. People resist this idea mightily, because as a justification for their own desires and opinions, they like to cite OBJECTIVE reasons for what they do -- that is, reasons grounded not in their own private understanding, but in some superior form of knowledge about the world as it actually is, knowledge that is not based on their own fallible perceptions and predictions but is TRUE.

When I say that all anyone can be concerned about and control is private perception, many people take this to imply an attack on the way they do things, and a recommendation that they behave differently. They interpret my words as if I had said they should stop being concerned with and trying to control other people and objective states of the environment, leave other people alone, and just be concerned with their own private lives. But this is not what I am saying at all. I am trying to tell them that even while they are trying to control other people and objective aspects of the environment, what they are really doing is controlling for their own perceptions. THEY NEVER HAVE BEEN ABLE TO CONTROL ANYTHING ELSE.

The resistance become mightiest from people who believe they are controlling other people for their own good. Not only do they insist that they must be doing good because that is what they intend, but they insist that the effects they have are OBJECTIVELY good for the other person.

In some ways what I really have to say is worse than telling people that they ought to stop controlling other people and be nice. At least they can fight back against such an attempt to tell them what to do. But I am discussing a description, not a prescription. I am saying that even when people think they are controlling other people, all they are actually controlling are their own perceptions. They can go right on doing what they're doing -- but it isn't what they think they're doing. I'm pointing out that this is the reason that they are so unsuccessful at controlling other people; they never were doing that in the first place, except in some trivial way that caused no problem for the other people. I am showing that when they have difficulties in achieving such apparent control of others, and try their best to overcome those difficulties, all they accomplish is to create conflict or put the other person in a state of reorganization that, in the end, preserves the other's capacities to control (or ends fatally).

The only time I use the term "ought" is in saying what people must do IF they want to avoid the difficulties. If you want to avoid conflict with others, then you have to stop trying to control what you can't control -- which is anything that matters to them, anything they are already controlling. If you like having those difficulties, if you think that conflict is an exciting and interesting state to be in relative to other people, then of course you needn't alter your ways. Evidently you are content with poor control of some of your perceptions or with imagining that you have good control when you don't, and if so that's your business (until you try it on me or someone I have decided to defend).

Before we can profitably get into a discussion how how people CAN interact under PCT, we must put aside all the misinterpretations of how they DO interact according to PCT. This means changing many informal interpretations of what we see going on around us.

Best,   Bill P.

Date: Wed Oct 21, 1992 11:48 am PST
Subject: Good work, Jim Graves

[From Bill Powers (921021.1130)]   Ed Ford (921021.0945) --

Ed, that was a great post about the AWOL problem. I am greatly impressed by your friend Jim Graves. I've always had difficulty in persuading therapists to "make a committment" to control theory without mixing it with remnants of contradictory theories. Jim seems to be giving HPCT a true test -- if it fails, there will be no doubt what theory is at fault.

David Goldstein's problems may not be exactly those that Jim Graves faces, in that (as I understand it) it wouldn't be legally possible to screen out potential residents who refuse to make the committment. But perhaps something equivalent could be done -- perhaps the screening could have to do with membership in one group or another, with the uncommitted residents undergoing a humane but basically custodial regime. To gain access to the other group, the committment of which Jim speaks would be required.

But I think that the intake procedure that Jim describes would itself be an enormous encouragement to make the committment. Involving the residents in the process is exactly what I would have recommended. Treat it as "our" problem, not "yours" or "mine." Get the conflict out of it as soon as possible.

I think we would all like to hear more from Jim. He and David Goldstein are running the first labs in social applications of control theory (other than your own).

Best,   Bill P.

Date: Thu Oct 22, 1992 5:01 am PST
Subject: Unimportant importance

From Greg Williams (921022)

WARNING: Hit your delete key now if you aren't interested in the PCT approach to social interactions -- this is liable to take several KB.

>Rick Marken (921021.1000)

>You can try all you want to control a control system but, as
>Powers points out eloquently and clearly in the Powers/Williams
>debate, your "control" is at best ephemeral (at least, when you
>are dealing with a control system that is organized the same as the
>control system that is trying to do the "behavior control"-- ie.
>a control system that controls the same perceptual world). When
>a control system tries to control other control systems the typical
>result is conflict -- unless you just want to see the control system
>produce an action that is irrelevant to the control system itself
>(the dog happily puts it's paw in the air to get all that dumb love
>that it really cares about).

"Ephemeral." Another new PCT-definition? The dog "ephemerally" raises its paw and walks beside its owner right into the neutering operating room. Look out, you critical reference signal, you're about to get Ace-of-Spayed!

>I think this is worth a discussion. I take the view that the way to
>promulgate PCT is to present the model and the research honestly,
>doing what we can to relate this to existing relevant concerns, but
>not shying away from explaining the true implications of the model.

So do I. I love that phrase: "the true implications of the model." Here, here!

>But I don't think it is worth it to compromise the model to try to get
>recruits -- PCT is neither a religion nor a political party.

Neither do I. And I don't think it is worth it to claim that PCT supports an ideology which it doesn't support, regardless of whether this gets or drives away recruits.

>It's not worth distorting the model to try to get people interested -- you
>just end up with people who are really interested in the version of PCT that
>you made up for their sake. But, I'd like to hear your point of view Greg.

Neither Bill nor I are attempting to distort the model itself. He and I differ to some degree on the possible details (particularly those which are hard to test at this time) of the model, but we differ most significantly on the importance to many people of some implications of the model WHICH WE BOTH AGREE ON. I claim that many people think that it is important to try to explain and deal with social interactions involving what you call "ephemeral control" and what Bill calls events which are "unimportant to the 'influencee.'" I don't deny that you can make such definitions as "ephemeral" and "unimportant"; I do deny that they have relevance to the many people who want to understand social interactions wherein parties are controlling their perceptions dependent on actions of other parties. The question of whether those people are misguided in some sense about what they think is important is an extra-PCT matter of ideological conflict. As I see it, PCT (undistorted!) has much to say about what these people think is important, even though you and Bill say that what they think is important involves ephemeral/unimportance. For these people (but not for you!), your ephemera/importance is beside the point.

>Bill Powers (921021.0915)

>I think you're still missing my point. The only way in which anyone
>can even know what actions he or she is producing is to perceive them.
>To "retrospectively consider having used those actions as being
>important" can mean, under PCT, only that the person experienced a
>perception dependent on the outputs at the time they were performed,
>remembered it, and later considered it to be an important perception
>(i.e., adopted a reference level for it).

I basically agree, with the exception that a person who does NOT remember his/her earlier actions can become convinced that he/she actually did them by receiving and accepting new information (such as a friend's explanation that "you signed the deed!" or a video showing the signing). But, so what? What is unimportant (your definition) at time x1 becomes important (your definition) at time x2. And, because this often occurs in the course of human life, lots of people are interested in situations where this is possible -- many of which are situations where one party is controlling his/her perceptions which depend on actions of another party, those actions being unimportant (your definition) to the second party at the time of the interaction, but important (your definition) to the second party at some time after the interaction. Note that the second party doesn't actually need to adopt a NEW reference level after the interaction: one can want to make money all along and think one is controlling for that, but find out (too late!) that he/she has actually lost money in the interaction. And that's one of the reasons for police and criminal courts showing GREAT interest in such interactions.

>If you can think of anything else beside a perception that can be controlled
>(or be important to a person) then you are proposing a different model from PCT.

I do not and am not, notwithstanding your own loose language about "controlling another's actions," which I have complained about before even as I went along with you on it.

>According to PCT, nothing CAN be of importance to a person but
>perceptions. This applies now, later, and to third-party observers.

I agree. Their CURRENT perceptions. At time t1, their perceptions then; At time t2, their perceptions then. At time t1, one's perceptions of one's actions occurring then can be unimportant (your definition), while at time t2 (>t1) the perceived memory of those actions can be important in the sense of causing a big problem for or making possible successful control of other perceptions (e.g., all their money is gone, or now they can rescue that drowning person).

>If you would try analyzing these situations in PCT terms, under which
>all the accountant or the little old lady can do is to control
>perceptions (including a perception of "gladness" and a perception
>that is "realized"), perhaps you would see that there is no
>contradiction. On the other hand, if you simply take appearances at
>face value and give them their traditional informal non-PCT naive
>realist interpretation, you will continue to miss my point.

I have not been doing and do not wish to do what your last sentence says. I don't want to be set up as a strawbuddy, either.

GW>>To be explicit, I claim that many people consider as highly
GW>>important social interactions of the type wherein B controls for
GW>>some of his/her perceptions which depend on some actions of A,
GW>>REGARDLESS of whether A's actions during the interaction are
GW>>important (your definition) to A or not ...

>Do they view those social interactions as important because they ARE
>being successfully controlled by the viewer, or because they are NOT
>being successfully controlled?

Neither. They view them as important because A later says that his/her actions which occurred during the interaction are (at that later time) important to him/her: "He tricked me into signing the deed." "It's a good thing I did those 10 laps of the pool each day like the teacher wanted, or I would have drowned, myself, out there!"

>Perhaps we can use the term "problem" or "difficulty" or some synonym
>to refer to perceptions that are important because attempts to control
>them do not work, and reserve the less specific term "important" to
>mean simply that the person perceives something, has a reference level
>for it, and acts to correct any difference (or would do so if
>possible). Thus to say that something is important to a person tells
>us that there is a reference level for a perception, but does not tell
>us whether the person is succeeding at controlling the perception. To
>say that a person has a difficulty with an important perception
>implies that attempts to control it are not working, and implies that
>reorganization is likely to be occurring.

Fine by me. Just don't ignore changes in what is a "problem" and what is "important" over time.

>And I think we should avoid further confusing the meanings of words by
>referring to THE importance of a social interaction or anything else,
>as if there were some objective standard of importance that is
>independent of anyone's perceptions or desires.

I continue to agree.

>I am not telling people what SHOULD be important to them. I am telling
>them that what IS important to them is their own perceptions. I am
>telling them that their perceptions are important because of what they
>desire those perceptions to be.

I have no problems with your definition, as I've said before.

>I am trying to tell them that even while they are trying to control
>other people and objective aspects of the environment, what they are
>really doing is controlling for their own perceptions. THEY NEVER HAVE
>BEEN ABLE TO CONTROL ANYTHING ELSE.

I'm saying the same thing, and trying to use PCT ideas to explain the nature and limits of controlling one's own perceptions which depend on others' actions.

>The resistance become mightiest from people who believe they are controlling
>other people for their own good. Not only do they insist that they must be
>doing good because that is what they intend, but they insist that the effects
>they have are OBJECTIVELY good for the other person.

The people I tend to respect are those who listen and (with reasonable caution) BELIEVE others when they claim that "what I did then is important to me now." Those who "objectively" disregard a "victim's" judgements about the importance to him/herself (the "victim") of others' controlling their (the others') perceptions depending on the victim's actions rate lowest in my own ideology.
I am saying that even when people think they are controlling other people, all they are actually controlling are their own perceptions.

I'm saying the same thing, and noting that often when a person controls his/her perceptions which depend on others' actions, the others and many third-parties (i.e., sociologists) think it important, either during or after the control episode.

>They can go right on doing what they're doing -- but it isn't what they
>think they're doing. I'm pointing out that this is the reason that they
>are so unsuccessful at controlling other people; they never were doing that
>in the first place, except in some trivial way that caused no problem for
>the other people. I am showing that when they have difficulties in achieving
>such apparent control of others, and try their best to overcome those
>difficulties, all they accomplish is to create conflict or put the other
>person in a state of reorganization that, in the end, preserves the other's
>capacities to control (or ends fatally).

But control of one's perceptions dependent on others' actions IS OFTEN SUCCESSFUL -- otherwise nobody would care about it! "Trivial" -- another new definition? Something can be perceived as "trivial" now and NOT "trivial" tomorrow. Why should anybody NOT be concerned about such phenomena? Because it is not PCT-control? Oh, come on! The FACT that anybody CANNOT make anybody else want what they don't want in the short-term is beside the point -- it looks to me like many people figured that out long ago, and went on to do what they CAN do (sometimes): control their own perceptions depending on others' actions.

>If you want to avoid conflict with others, then you have to stop trying to
>control what you can't control -- which is anything that matters to them,
>anything they are already controlling.

I agree. I think I see so little conflict in my own everyday life (speaking, of course, from the distinctly privileged viewpoint of a farm in central Kentucky!) because the "trivial" type of control (your word), as exercised both by myself and my acquaintances, is so often successful. When the "trivial" type of control is occasionally UNSuccessful, I begin to appreciate its importance ever more!

Best,   Greg

Date: Thu Oct 22, 1992 9:12 am PST
Subject: Rick's paper

[From: Bruce Nevin (Thu 921022 10:46:41)]

Rick, I just read through your behavior of perception paper this morning. A few minor comments and questions.

Under "A Perceptual Control Hierarchy":

  The output transducer amplifies and converts this difference into
  actions which affect the environment or [become ==> into] reference
  signals for lower level systems

At the end of the section "Hierarchical Invariance":

>   Most hierarchical models of behavior require that a high level
>command be decomposed into the many lower level commands that produce
>the intended result. In the hierarchical control model, both the high level
>command and the intended result of the command are represented by a
>single, unidimensional signal. The signal that represents the intended result
>is a function of results produced by many lower level commands. But the

>high level command does not need to be decomposed into all the appropriate
>lower level commands (Powers, 1979). The difference between the high
>level command and the perceptual result of that command is sufficient to
>produce the lower level commands that keep the perceptual result at the
>commanded value (Marken, 1990).

The last sentence is unnecessarily obscure, I think. Maybe it would help to state that the high-level "command" is nothing other than the reference signal r identified near the outset of the paper. After all, the term "command" is used here ironically at best, no?

Last sentence before "Levels of Perception":

  . . . has time to control the relationship between mouse [movement]
  and cursor movement.

Last sentence before "The Relationship Between Behavior and Perception", vs. second sentence after that heading: are phonemes sound configurations or sound events? I think maybe the difficulty lies in the fact that events are sequences too, only they are short and they are very well learned (automatized). How about:

  An actor can produce a desired sequence of sounds, for example, by
  speaking a word. (A word is a sound event comprising a short sequence of
  phonemes whose order is so well learned as to be automatised.)

However, this may be an unwanted complexity at the particular place that it occurs in the paper. I'm not sure what resolution you would prefer in that case, except that presumably you don't want to say that a phoneme is an event.

References: No date for the Albus reference.

Question: what results does PCT predict that standard theories would not, or what results do standard theories predict that PCT does not? You show that phenomena treated separately in standard theories (viz. hierarchical structure of behavior and of perception) are unified in PCT. What are the consequences for standard theories of treating them separately, which may be avoided by PCT? Why should they care? I am reminded of the bank robber who spills some of the loot and puts his gun on the counter so he can pick it up, and when he turns to make his escape the teller (with hands up) says "you forgot something."

I appreciated receiving this, and it has ticked over some mostly subconscious processes that I hope will result eventually in a paper relevant to linguistics.

        Be well,   Bruce   bn@bbn.com

Date: Thu Oct 22, 1992 11:59 am PST
Subject: Unimportant importance

[From Rick Marken (921022.1000)]   Greg Williams (921022)--

>And I don't think it is worth it to claim that PCT supports an
>ideology which it doesn't support, regardless of whether this
>gets or drives away recruits.

Is it the "people cannot be controlled" ideology? I think talking about this does cause a lot of problems, especially if we are not clear about terms. When you actually work with the working model you can see what you can and cannot do to it. If you want to call some of those things "control" that's fine. If, as Bill said, you are happy with the results of your interactions with control systems, then that's great -- whether you want to call what you do "controlling" or "educating" or "cooperating" or whatever. The control model just happens to work the way it works. If you try to control some variable aspect of the control system's performance then you will be successful if it is something that is not also being controlled by the control system or you will get into a conflict with the control system if it is. If you are "controlling" and not getting into conflict then either you are controlling what the control system is not controlling or (most likely) you are not really controlling (bringing a perception to a preselected reference level and maintaining it there against disturbance). If people are control systems, then this is just the way it works.

[On re-reading this I see that I WENT UP A LEVEL right at this point. I thank Greg's tenaciousness for this consciousness raising.]

I admit that, in my discussions of PCT, I have revealed my personal ideology -- which is to avoid conflict (especially the violent type) and foster cooperative efforts to control mutually controlled variables. I may have given the impression that I think PCT justifies this ideology -- IT DOES NOT. Maybe this is what Greg is getting at in his critique of PCT ideology. If so, I repent and accept your criticism. PCT only says that controlling other control systems (REALLY controlling them; not the mamby pamby stuff) leads to conflict -- IF you try to control what the other control system is also trying to control. The ideology part is thinking that this kind of conflict is no good; ie. having a reference signal set at 0 for conflict. If you like conflict (and many people seem to love it -- football games, free enterprise economics, etc) then PCT can, indeed, help you produce all you want.

It is hard to discuss PCT without letting my references for non-PCT perceptions get in the way. I see that I have been guilty of this -- Bill's last post really made it clear to me. It is incredible how hard it it to describe the model without biasing one's description in terms of one's principles. I want to see a world where people are not at each other's throats; but that has nothing to do with PCT, except that PCT can help show people how people might be able to live in such a world (just stop trying to control each other). But the goal of living in such a world is mine; others might like to live in a world of hand to hand combat and cut-throat competition. Different reference levels for the same principle. I guess I'll just have to take a deep breath and accept the fact that other people may really WANT to live in a world like we live in -- filled with hatred and oppresion (ie. conflict). Fine with me -- except (as Bill said) when they try to oppress me or my loved ones.

So, let's get back to the model, knowing that it is difficult to DISCUSS it without coloring the "implications" with one's own values. That's why I like discussing the model in terms of the computer demos -- we can get closer to seeing what the model DOES, not what we think it might IMPLY about other stuff we care about.

Best regards   Rick

Date: Thu Oct 22, 1992 12:18 pm PST
Subject: awols

David Goldstein 10/21/92

Thanks for the suggestions. On second thought, the real reference level is zero residents killed when they are at our center. When residents go AWOL, there is an increased safety risk which they do not perceive. AWOLing is just one way in which a resident can be unsafe.

Here are my reactions to your suggestions:

(1) get a committment during the pre-admission process--One of the areas which lead to residential placement is what we call "Promptness and Attendance." It includes things such as: does not report to expected destination, does not remain at assigned destination, school truancy, general runaway. If this is one of the areas which require residential placement, it is unrealistic to expect a miracle cure based on a verbal committment. In spite of what I just said, it is worth trying.

(2) motivate the residents to want to stay in the program by making it interesting and attractive--We already do this. The program has to be changed often because residents become bored easily. Even a program which is effective for some residents may not work with all residents. It is obvious that any program can be improved and made more interesting and attractive.

I think that Bill and Rick are correct. There is no way to guarantee that a resident will not go AWOL except if a resident wants this. All the different things we are doing or could do cannot 100% ensure that this will be a goal for a resident. There is no foolproof way of knowing whether this is a goal for a resident.

Some final thoughts: (a) if a resident continues to go AWOL and cannot be persuaded to stop this, the safest thing is to dicharge the resident with the recommendation of a more restrictive setting, (b) a set of items could be developed which could be qsorted to describe a resident's view about AWOLS, (c) after an AWOL, the therapist should conduct a post-critical incident counseling session to analyze the reasons for it and come up with some recommendations to make it not necessary if possible.

The general issue of control versus influence is involved in the AWOL example. Only a resident can control his/her AWOLS. Staff can try to influence a resident but cannot control the AWOLS without taking measures which are not acceptable.

Date: Thu Oct 22, 1992 1:55 pm PST
Subject: Importance out, problems in

[From Bill Powers (921022.0900)]   Greg Williams (921022) --

>I claim that many people think that it is important to try to
>explain and deal with social interactions involving what you call
>"ephemeral control" and what Bill calls events which are
>"unimportant to the 'influencee.'"

Formerly I was thinking that "important" meant "causing some sort of problem that had to be dealt with." Now I am saying that a perception is important to a person if the person perceives and has a reference level for it and if possible controls it. Now we can say that unimportant variables are those that are not perceived or that are perceived but have no preferred state, while important perceptions are those that people are actively concerned about controlling. So now we agree, if you accept this new definition, that all forms of control are important to the controller, whether they involve the actions of other people or not. This simply makes "important" synonymous with "controlled or potentially controllable." Because we are not concerned with uncontrolled perceptions right now, we can drop the term "importance."

Let's see how far you will go along with this development.

A "disturbance" is either a direct influence applied to a controlled variable, or a change in the parameters of the link between a person's action and the variable it controls.

A "problem" is a situation in which control is sufficiently difficult to result in substantial and sustained deviation of one or more perceptions from their reference levels.

"Reorganization" is used in a very general sense, and could include a change in the operation of higher-level systems without any actual change in organization at those levels (i.e., a change in strategy dictated by learned principles). Sorting out the ways in which behavior at a given level can change its characteristics can be left for later.

First, let's consider a non-exhaustive set of cases from the standpoint (mainly) of the person whose behavior is affected from outside.

1. Unproblematic interactions:

If your actions (outputs) are controlled by another person, but in such a way that NO perception controlled by you is materially disturbed in the process (meaning that the other's disturbance plus your action keeps your perception near its reference level), then the other has not caused a present-time problem for you. It makes no difference to you whether the other elicited an action intentionally or unintentionally, because in neither case do you experience a error that can't easily be counteracted. You don't need to distinguish between disturbances applied for a purpose and disturbances that occur naturally or accidentally.

I claim that this is the most common form of human interaction: all parties involved continue to control their own perceptions without any problems, even though they continually adjust their actions to compensate for disturbances by the actions of other people, and often deliberately elicit actions from other people (handing the cashier your $5 purchase and a $20 bill).

2. Present-time problem-causing interactions:

Disturbances of your controlled perceptions will cause a problem in present time when either the disturbance exceeds your capacity to resist it, or the action necessary to resist it has side-effects that disturb another of your perceptions in a way that can't be resisted.

Direct disturbances that cause errors result in conflict between persons. Disturbances that elicit actions which cause errors in the actor result in internal conflict in the affected person.

When a present-time problem is caused by an interaction, the person experiencing the uncorrectable error can only endure the error or reorganize.

3. Delayed problem-causing interactions:

A disturbance can result in an opposing action that has problematical effects which are not immediately experienced. If you anticipate those effects (all anticipations or predictions occur in present time), an internal conflict will immediately result. If you carry out the action that resists the present-time disturbance, you will cause a departure of a predicted future state of a perception from the desired future state. This is a present-time error. So this is really case 2, above.

If you fail to anticipate or incorrectly anticipate the future effect of your action, then you will do nothing to prevent its occurrance. You will produce the action needed to counter the present-time disturbance and you will continue to control successfully until the delayed effect occurs. The problem will then appear, in what is now present time. If you can resist the unwanted effect, there will be no problem. If you can't resist the effect, you will either suffer the resulting error, or reorganize.

Now let's shift the point of view more toward the person who is controlling the actions of someone else.

"Output" means the physical output generated by a control system.

"Action" means a perception of an output (either one's own or someone else's).

4. Unproblematic control of the actions of others.

Present-time control:

Applying disturbances that can be resisted by another person can be used by a controller to control the outputs of the other person. The aspect of those outputs that is controlled is whatever aspect is perceived by the controller as an action. If the perceived action corresponds exactly to the aspect of output that affects the other's controlled variable, the other's physical output is controlled when the perceived action is controlled.

Future control:

Disturbances can be applied in a way that elicits an action that entails a predicted future effect on the other person. If no uncorrectable error is caused in the other person, either present or anticipated, the action will take place and the controller will immediately experience a match of the predicted future effect to the effect the controller wants. It does not matter whether the effect will actually occur in the future, because the prediction is made in present time and the goal is satisfied in present time. If, when the future arrives, the effect does occur, the controller will continue to experience zero error. If this future effect does not cause any error in the controlled person, the controlled person also will experience no error.

5. Problematic control of the actions of others.

Present-time problems.

If a disturbance materially alters a controlled variable in another person, two things will happen. First, the relationship between the disturbance and the action it controls will change. If that change is large enough for the other to lose control, control of the other's action will be lost. Second, the other person will begin to reorganize. That will alter the characteristics of the control system in question, and may also bring other control processes into play aimed at correcting the error. In any case, conflict between the controller and the controlled person will appear.

If the controller has to produce too much effort to maintain control, or loses control altogether, the controller will begin to experience uncorrectable error and will reorganize.

Future problems.

If the controlled person anticipates a future error as a result of a present action, the action will not take place as the controller wishes. The controller will experience error of two kinds: the present-time action will not occur as desired, and the prediction of the future effects of the action will be different from the effect that is wanted in the future. The controller will either continue to experience error or reorganize.

If the controlled person does not anticipate the future error, the action desired will take place and the controller will be satisfied. When the future arrives, however, the effect on the controlled person (if it occurs as predicted) may cause an error. The controlled person will then act to oppose the effect. If this action by the controlled person is successful, the controller will experience an error and reorganize. If it is unsuccessful, the controlled person will experience an error and reorganize. In either case, one of the persons must experience an error, and interpersonal conflict will exist.

In summary:

Unproblematic control of another person's actions causes no uncorrectable errors in either the controller or the controlled person. Each person continues to operate normally, without any change in organization. Each prevents the other from having any unwanted effect on any controlled perception. All parties adjust their outputs as required to maintain control.

Problematic control of one person by another results in conflict between the parties, and loss of control by one or both of them.

OK so far?

Best, Bill P.

Date: Thu Oct 22, 1992 2:15 pm PST
Subject: PCT linguistics; Committment

[From Bill Powers (921022.1200)] Bruce Nevin (921022.1046) --

I'm really pleased to hear that you're thinking of writing a paper on PCT and linguistics.

I've been meaning to ask -- in your latest foray into Achumawi (sp?) land, did you have any time to check out PCT concepts across cultures?

David Goldstein (921022) --

The AWOL problem, I think, is a symptom, not something to be treated in itself. If you deal with the right problems, the AWOL problem will go away by itself (as much as possible).

>[The reason for placement]... includes things such as: does not
>report to expected destination, does not remain at assigned
>destination, school truancy, general runaway. If this is one of the
>areas which require residential placement, it is unrealistic to
>expect a miracle cure based on a verbal committment.

In "getting a committment" I don't think that the committment you want is simply not to do the things that a complainant complained about. That's just confrontational. It would be like saying "Promise you'll stop doing those bad things." I agree that this would be futile.

Probably Ed Ford or Jim Graves could put this better than I can. I would guess that the kind of committment you want is to learn how to get along better with people, be happier, feel confident, understand more, avoid things (like being in this institution) that you really don't want to happen, have some hope in the future. In other words, get a committment to make a real try at getting things the person really wants, and to allow others to help where it's needed.

I don't think that not remaining at an assigned destination is a particularly horrible crime, and the resident probably doesn't, either. The crime, from the description, seems to be not letting others have total control over you. I should think you would agree with the residents that this is a bum rap. But bum rap or not, there are many good things that can be done as long as the resident has to go through this -- learning how to survive such a coercive system without getting in trouble, for example, or even coming out of this a lot happier than before about a lot of things. That's what you need a committment for. You just want the resident to commit to having a real try at a better life.

Best to all,   Bill P.

Date: Thu Oct 22, 1992 3:31 pm PST
Subject: closed loop

from Ed Ford (921022.1330)

Because of house building by Greg and traveling by Ed, Closed Loop will be mailed sometime early the week of Nov. 2nd.

Greg, don't forget the addresses you were going to send.

Best, Ed

Date: Thu Oct 22, 1992 4:46 pm PST
Subject: Re: AWOLS, PCT Popularity

[From Dick Robertson 921022]

Rick Marken advocates doing our thing and not wasting anymore time trying to "tell me what to do , don't waste my time telling me abstract stuff about how usefull for eventually helping neurosurgeons figuring out where to cut and so or" in the 1960 paper and Rick you have built nicely on that in your more recen So I vote for more experimentation and modeling instead of trying to convert the

Best, Dick

Date: Thu Oct 22, 1992 6:21 pm PST
Subject: Re: PCT popularity; Why 99%?

Re: PCT Popularity; and Why 99%?   From Tom Bourbon (921022.13:45)

I second Rick Marken's claim that we should not neuter PCT in an attempt to gain a wider audience. I will describe some of the pressures I have encountered to do that. Then I will identify several people who have published extensively, claiming that control theory is compatible with everything else in behavioral- social-cognitive science. Then I will describe a clear example of the disappointing consequences that follow from their having published distorted, simpatico versions of control theory.

[Rick Marken (921021.1000)]

>>Greg Williams says:

>>Bill, who am I to tell you that you should perceive a problem
>>with PCT being a tough sell? That I believe there IS a problem
>>with it is MY ideology. There is no perceived problem for you if
>>you believe that you are right about what everyone SHOULD believe
>>to be important, and that virtually everyone else is wrong, even
>>if you believe PCT ideas deserve wider attention... until you

>>begin to perceive people shrugging and walking away when you tell
>>them what they should and shouldn't believe to be important. And
>>then I won't have to tell you that you should perceive a problem;
>>you'll be perceiving it.

Rick Marken replies:

>I think this gets at the heart of Greg's complaint; he would
>really like to see PCT ideas get wider attention. I think he sees
>a lack of interest on the part of some PCTers (like myself) in
>finding common ground with psychologists, roboticists, biologists,
>AIers, ALifers, etc etc -- ie. with others in the community of
>life scientists who might profit from an examination of PCT ideas.

##############
TB:

Rick describes a frequent interpretation offered by reviewers, editors, and others who see or hear manuscripts and presentations on PCT, especially when the presentations or manuscripts come from "hardcore" PCTers. I have a collection of reviews in which the writers say I (and my co-authors if there are any) went out of my way to make PCT unpalatable, or that I want readers to reject PCT. Most of them continue with remarks that I should point out how PCT "is like --- ;" or "is similar to --- ;" or "is just another way of saying --- ." Or THEY say I should say that "we (they) already know --- ." A few say that most strong assertions by PCT writers (eg., behavior controls perception) are "merely ideological." etc., etc., etc.. When we decline to mend our evil ways (by going along with what the reviewers say) our refusal is often taken as proof that we do not want people to read about or to understand PCT. Of course, their comments are self-fulfilling: after they tell the editors not to publish our manuscripts, no one reads about PCT.

I believe it is essential that we avoid presenting PCT in a watered-down version and that we resist all suggestions that it "offers another perspective" on the same old things, or that it is "a convenient framework" for "viewing" and unifying things "we already know." Below, I will present a clear example of why I think we must hew close to the basics.

#######################

RM:

>Those of us who do not seem to "compromise", "see commonalities",
>etc believe that we are just presenting the PCT model -- we don't
>feel that there is an agenda to alienate potential friends; but
>apparently it seems like this is true to some people (like Greg,
>I think).

##################
TB:

Rick's point is well taken. You have not lived life to the fullest until you labor to submit a manuscript on PCT (perhaps for the second or third or ... time) then read a reviewer's smug accusation that you (the one who wrote the manuscript) are "determined to elicit rejection." It calms the spirit and soothes the nerves.

####################
RM:

>I think this is worth a discussion. I take the view that the way
>to promulgate PCT is to present the model and the research
>honestly, doing what we can to relate this to existing relevant
>concerns, but not shying away from explaining the true
>implications of the model. I believe that attempts to "find common
>ground" produce the Carver/Scheier approach to PCT -- which ends
>up using the terminology of PCT but misses the basic point (and
>succumbs to the causal view of behavior in the end). I realize
>that the approach I advocate is not a good way to drum up a PCT
>following -- most people do shrug and walk away simply because
>they don't see the problem that PCT solves -- or don't get PCT

>even if they do see the problem. But I don't think it is worth it
>to compromise the model to try to get recruits -- PCT is neither
>a religion nor a political party. And every so often someone DOES
>get stoked on it (like I did). I think we are gaining PCTers who
>really get it (ie. model-ers) at the rate of about 1 a year now.
>That's plenty for me. I just don't think you can MAKE people be
>interested in PCT. When the light does go on in a person, PCT
>sells itself. It's not worth distorting the model to try to get
>people interested -- you just end up with people who are really
>interested in the version of PCT that you made up for their sake.
>But, I'd like to hear your point of view Greg.

####################
TB:

Agreed, on practically every count. It is a mistake to distort, water down, neutralize, or defang PCT. We should not go out of our way to build bridges, identify communalities, find common ground, etc, when the other side of the river is quicksand. Control BY an organism is different from control OF an organism (even if "control of" is camouflaged in the contemporary jargon of "behavioral analysis" or "cognitive science"). Period. Crisp predictions by a generative model of control behavior are not the same as statistical mush, in which significant differences between mean scores from groups, and low but "significant" correlations, are offered as evidence supporting one or another "theory" of behavior. Period. Usually the behavioral phenomena or the cognitive- emotional-social processes alleged to exist on the other side of the river are phantoms.

In such cases, there is nothing to which we can build a bridge -- there is no communality. But that does not prevent some writers from publishing extensively with claims that PCT can be all things to all people. Rick mentioned Carver and Scheier. They are part of the list of "villains" that rolls from my tongue or fingertips as one entity:

Carver-Scheier-Hyland-Lord-Hollenbeck.

That would be C.S. Carver, M.F. Scheier, M. Hyland, R.G. Lord, and J.R. Hollenbeck. There are others, but this group deserves special attention. All offer CT as a "framework" or "perspective" or "conceptualization" or "view" for virtually everything. All "build bridges," "promote unity," "integrate," and all of the other things PCTers are so often urged to do. Collectively, the members of my rogues' gallery have published tens of times more material on what they call control theory than have any participants on CSG-L. For them, publication is easy -- say that control theory is not a threat, that it affords another compatible perspective on everything, and no one is bothered. Isn't that the best way to spread the word about PCT?

No!

Their presentations of a nonfunctional and eviscerated control theory have done far more harm than good. On this opinion, I cannot be moved. Many people have formed their "understanding" of control theory by reading the numerous publications of that group. Their collective writings are so extensive that an innocent reader could easily believe they form an authoritative literature. That is not true. Time for my major case in point.

Bandura, Albert (1989). Human agency in social cognitive theory. American Psychologist, 44, 1175-1184.

As part of his presentation, Bandura raised and knocked down a pathetic description of a "negative-feedback" system. That exercise occupies much of pages 1179-1191. To any PCTer, it was obvious that he did not understand negative feedback, but Bandura is an authority and now the article is cited widely and favorably.

Bill Powers wrote a "comment:"

Powers, William T. (1991). Commentary on Bandura's "Human agency." American Psychologist, 46, 151-153. I recommend it.

A few others also submitted comments, some dripping with praise. Of course, Bandura wrote a reply:

Bandura, Albert (1991). Human agency: The rhetoric and the reality. American Psychologist, 46, 157-162.

Bandura missed the point of Bill's comments and continued to discuss control theory in thoroughly negative tones. While I was reading his reply, the awareness dawned that he was writing about the control theory presented by the unholy alliance of Carver-Scheier-Hyland-Lord-Hollenbeck, not about PCT. My insight was confirmed when I read:

"Locke (in press) has argued that much of control theory involves translation of the principles and knowledge of goal theory into stilted machine language without providing a new perspective or predictive benefits. He further showed that adherents of control theory have now grafted so many ideas from other theories on the negative feedback loop as (sic) to remedy its prediction problems that control theory has lost its distinctiveness" (1991, p. 158).

Sadly, those are exactly my own conclusions when I assess the literature of the devil's alliance. I knew Locke would cite the popularizers and bridge builders.

He did:

Locke, Edwin A. (1991). Goal theory vs. control theory: Contrasting approaches to understanding work motivation. Motivation and Emotion, 15, 9-28.

Locke accurately summarized much of the material produced by the "nice guys" -- the sweetness and light brigade -- of control theory. EVERYONE who believes we should go out of our way to mollify people from every other camp in cognitive-social-behavioral science should read that article.

Weak-kneed presentations of PCT, in which core concepts are abandoned or verbally "modified" every time some established critic squeaks, do no good. They are misleading and harmful.

I am not suggesting that we go to the other extreme and bash everyone who does not catch on to PCT as quickly as we might like. But anyone who believes we will spread awareness of PCT more quickly by deliberately softening its implications should read the references I have cited, in chronological order.

Why 99%?

Earlier this year, a former graduate student and I submitted a manuscript describing our modeling of cooperation by pairs of people. Our results include numerous +.997 correlations between predictions of moment-by-moment actions made by two interacting PCT models and the actions of two people. In the manuscript, we cited the exchange between Bandura and Powers, in American Psychologist, and tried to build a bridge -- a real one, not a string of b*** s***. We suggested that Bandura's misunderstanding of control theory and of negative feedback came from his familiarity with faulty sources -- the nasty five. We also suggested that our results provided a modest example of the predictive power of a legitimate negative feedback model, namely, the model from PCT. We were rejected. Among the many fascinating reasons, one reviewer "assured the authors" that "Bandura would not be impressed." (Who knows, maybe that review was by Bandura himself!)

Assuming the reviewer was right, what WOULD impress Bandura? He gives a strong hint in his reply in American. Psychologist:

"As shown in Table 1, perceived self-efficacy accounts for a substantial amount of variance in phobic behavior when anticipated anxiety is partialed out, whereas the relationship between anticipated anxiety and phobic behavior essentially disappears when perceived self-efficacy is partialed out" (1991, p. 160).

Forget about the problems of defining terms and constructs -- quicksand and phantoms, all! Table 1 is on page 161. In it is a summary of correlations from several studies on the aforementioned phantoms. They range from -.22 to +.77. They are accompanied by a cloud of asterisks: * ** ***. How silly I was not to see why Bandura would be unimpressed by correlations of +.997! Our correlations "accounted for" 99.4% of the variance; he must only be impressed if you can "account for" a paltry 4.8% to 59%. Our coefficient of alienation (aka, probability of failure in a prediction) would be < 0.1;

if his representative, the reviewer, is correct, Bandura must want probabilities of failure ranging from > 0.98 to about 0.66. Now I understand.

(The "coefficient of alienation," or "coefficient of failure," was a big topic on CSG-L long ago. For a good source -- an old one -- see:

Guilford, J.P. (1956). Fundamental statistics in psychology and education. NY: McGraw-Hill.)

Why strive for 99% of the variance "accounted for?" It is simple: To avoid the mistake of believing phantoms are facts; to avoid the sham and scam of saying you have a science, when all you have is statistical mush.

Why insist on hewing as close to the core of PCT as possible, rather than making PCT an easy pill for all to swallow? Figure it out.

Best wishes,

Tom Bourbon
MEG Laboratory
1528 Postoffice Street
Galveston, TX 77550
USA

e-mail in care of:
PAPANICOLAOU@UTMBEACH.BITNET
PAPANICOLAOU@BEACH.UTMB.EDU
PHONE (409) 763-6325
FAX (409) 762-9961

Date: Thu Oct 22, 1992 6:51 pm PST
Greg Williams (921022), Rick Marken (921022.1000), Bill Powers (921022.0900)

A lovely sequence today. More eloquent clarification, going up a level and expansion of control over time.

I think time is an important variable here, which deserves recognition and will clarify a number of the concerns we all wrestle with in real time. Time figures prominently in the discussion of AWOL commitment as well. Some thoughts on time:

The hierarchical control mathematics taught by Bill and Rick includes slowing factors and thus time recognition. I prefer to think and teach graphically. In Durango 1991, I presented a chart in three levels, portraying what I called: Timing of control. Two levels are shown here, which is all we need.

```
                     -------------------------->>>>  TIME

Level 2  |              | r           |
Slow!    |    __________v______________________________________
         |   |                                                 |
         o-->|                    comparator                   |---o
       p |   |_________________________________________________|   | e
    _________|_______     _________________________________________v_____
   |                 |   |                                               |
   | sensor / input  |   |              amplifier / output               |
   |_________________|   |_______________________________________________|
         ^                    |                         |
         |                    |       |                 |
         |                    |       |                 |
Level 1  |          o---------o       |                 |
Faster!  |          | r t1            |                 | r t1 (unchanged)
         |    ______v_____    TIME    |  TIME      _____v______
         |   |            |     t1    |   t2      |            |
         o-->| comparator |---o       |       o-->| comparator |---o
     p t1|   |____________|   | e t1  |  p t2 |   |____________|   | e t2
    _____|_____          _____v_____  |  _____|_____          _____v_____
   |           |        |           | | |           |        |           |
   |  sensor   |        | amplifier | | |  sensor   |        | amplifier |
===|___________|========|___________|=|=|___________|========|___________|===
         ^                    |       |       ^                    |
    _____|_____          _____v_____  |  _____|_____          _____v_____
   | controlled|        | amplified | | | controlled|        | amplified |
   |  input    |        |  output   | | |  input    |        |  output   |
   | variable  |        | variable  | | | variable  |        | variable  |
   |   i t1    |<-------|   o t1    | | |   i t2    |<-------|   o t2    |
   |___________|        |___________| | |___________|        |___________|
         ^                            |       ^
    _____|_____                       |  _____|_____
   |  disturb- |                      | |  disturb- |
   |   ance    |                      | |   ance    |
   |   d t1    |                      | |   d t2    |
   |___________|                      | |___________|
```

This is an attempt to portray that the higher level control system MUST be slower for the combined system to be stable. We demo this easily with hand movements.

I also wrestled with time when I developed and posted on behavior of perception (Dag Forssell (920926)). The intensity level of control - muscle fiber control - is here and now. The configuration level - position of body part is rather present also. The sequence level - driving in progress - covers minutes or hours. The systems concept level - I am a professional - covers almost infinite time.

It feels awkward to portray the hierarchy of control without all the levels being focused on the present. In today's post Bill clarifies that we consider a large range of time in our imagination. This is helpful to me. It occurs to me that the time aspect of HPCT can be portrayed as follows: (In my world, if I cannot graph it, it is not real |:)   )

```
                     ------------------------------>>> TIME

Syst Conc ***************************************************************
Principle                **************************
Program                       *****************
Sequence                        *************
Category                           *******
Relationship                        *****
Event                                ***
Configuration                         *
Transition                            *
Sensation                             *
Intensity                             *
```

I have continued to work on behavior of perception, and will send the graphs to any netter who asks politely with snail mail address.

Wayne Hershberger:

Are you interested in a set of blank forms you can use to portray the vision processes. How many levels of perception? Did you keep (920926) posting on that? I think it would be most interesting to work out a good set of charts on vision. By the way, National Geographic for November has a feature article on vision: THE SENSE OF SIGHT, complete with a spread: PATHWAYS TO PERCEPTION that shows details of THE RETINA, a sense of hierarchy in PUTTING IT TOGETHER. I wonder how much of these colorful illustrations hold water.

Gary Cziko, Greg Williams & Bill Silvert:

I'll hold the revised Starter Document a few days more, hoping to resolve the downloading of binary files to general satisfaction with the help and ingenuity of Pat and Greg. Greg thinks Pat can create an ASCII file for Bill's server which can be downloaded through regular mail. With instructions included at the beginning of the file for the use of DOS Debug program, this file is converted to uud.exe, the dos enhanced decoder. Did I get that right, Greg? Bill, may I state in the starter.doc that the file is on the server as shown:

```
programs/source:        ASCII files  (Bill, are they)?

uud.c            13750  Enhanced decoder, C source
uue.c             4734  Enhanced encoder, C source
uudself.dos     ??????  Compile uud.exe with DOS debug. Directions included.
```

Best to all,   Dag

Date: Fri Oct 23, 1992 1:27 am PST
Subject: Re: Rick's paper

[From Marcos Rodrigues 921023.1000]

Bruce Nevin: Your message of Thu, 22 Oct 92 11:25:09 EDT:

>The last sentence is unnecessarily obscure, I think. Maybe it would
>help to state that the high-level "command" is nothing other than the
>reference signal r identified near the outset of the paper. After all,
>the term "command" is used here ironically at best, no?

Having read the extract you posted, I also think it is obscure. However, I think that the word "command" is strictly right, since the reference signal is a request for a perception.

Regards,   Marcos.

Date: Fri Oct 23, 1992 3:44 am PST
From: Hortideas Publishing / MCI ID: 497-2767

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Need uuencode, too

From Greg Williams (921023 - direct)

Hi Dag,

Please send uue.c also, so we can check to make sure our uud.exe works. No, we don't have it done yet -- good weather all this week, so housebuilding is taking priority. Probably this weekend. Thanks for being patient.

Nice comments on the recent me/Rick/Bill comments. Maybe we can all be happy with the outcome if nobody thinks they had to "give up" very much.

Best,   Greg

Date: Fri Oct 23, 1992 4:01 am PST
Subject: Ideologies; substantial agreement with Bill

From Greg Williams (921023)   >Rick Marken (921022.1000)

>Is it the "people cannot be controlled" ideology? I think talking about
>this does cause a lot of problems, especially if we are not clear
>about terms. When you actually work with the working model you can
>see what you can and cannot do to it. If you want to call some of
>those things "control" that's fine. If, as Bill said, you are happy
>with the results of your interactions with control systems, then that's
>great -- whether you want to call what you do "controlling" or
>"educating" or "cooperating" or whatever. The control model just
>happens to work the way it works. If you try to control some variable
>aspect of the control system's performance then you will be successful
>if it is something that is not also being controlled by the control
>system or you will get into a conflict with the control system if it
>is. If you are "controlling" and not getting into conflict then either
>you are controlling what the control system is not controlling or
>(most likely) you are not really controlling (bringing a perception
>to a preselected reference level and maintaining it there against
>disturbance). If people are control systems, then this is just the
>way it works.

One of the ideologies I've been claiming is unsupported by PCT is that which gives great significance to the fact that "if you try to control [your perception of -- I hasten to add!] some variable aspect of the control system's performance then... you will get into conflict with the control system if it is [something that is also being controlled by the control system]" but gives minimal significance to the fact that "if you try to control [your perception of -- I hasten to add!] some variable aspect of the control system's performance then you will be successful [sometimes -- I hasten to add!] if it is something that is not also being controlled by the control system...." I claim that many people are interested in instances of the latter, and that those people will perceive PCT as irrelevant to their interests if PCTers make a big deal of the former fact and minimize the significance of the latter fact. Many people aren't bothered if they can't make you want something you don't want, because they don't care about CHANGING what you want -- they are satisfied to be able to control their own perceptions which depend on (some of) your actions by having a sufficiently accurate model of (some of) your current wants. And many people (including, I suppose, you) are interested in instances of the latter because of subsequent problems or problems solved FOR THEM, IN THEIR OPINIONS, after their having participated in the instances on the "influencee" end. And other people (third-party investigators) are interested in instances of the latter, also, because such instances often are judged as important by some or all parties involved in them, at least retrospectively.

All along I have been trying to be careful in saying what one can control, namely, some of one's own perceptions (including some which depend on others' actions). Bill has been speaking of "controlling others' actions," but I don't applaud that loose phrasing.

>I admit that, in my discussions of PCT, I have revealed my personal
>ideology -- which is to avoid conflict (especially the violent type)
>and foster cooperative efforts to control mutually controlled variables.
>I may have given the impression that I think PCT justifies this
>ideology -- IT DOES NOT. Maybe this is what Greg is getting at in
>his critique of PCT ideology. If so, I repent and accept your criticism.

No, it isn't. I doubt that nonPCTers will be turned off by that ideology; many of them also believe that avoiding conflict is important, and they are looking for ways to achieve that aim.

>PCT only says that controlling other control systems (REALLY controlling
>them; not the mamby pamby stuff) leads to conflict -- IF you try to
>control what the other control system is also trying to control.

I claim that "real control" includes controlling (some of) your perceptions which depend on (some of) the other control system's actions. But you imply that it means controlling (some of) your perceptions of (some of) the other control system's controlled perceptions... or trying to do so, since this leads to conflict. I think Bill agrees with me on this, now. In his most recent post, Bill defines "A's action" as "A's output as perceived by B," so he now appears to accept my definition of "real" control of one's perceptions which depend on others' actions.

>So, let's get back to the model, knowing that it is difficult to
>DISCUSS it without coloring the "implications" with one's own values.

Three cheers for that!

>Bill Powers (921022.0900)

>Let's see how far you will go along with this development.

So far, I almost completely agree with what you say (see minor quibbles below). Now, what about future-time facilitations, as well as problems? I suppose we are still in disagreement about that. Also, I would like to see you also develop similar schemas for third-parties who are interested in such interactions.

>If the controlled person does not anticipate the future error, the
>action desired will take place and the controller will be satisfied.
>When the future arrives, however, the effect on the controlled person
>(if it occurs as predicted) may cause an error. The controlled person
>will then act to oppose the effect.

IF such action is possible. There might be no way to "right the wrong" -- no successful (or maybe even partly successful) action possible.

>If this action by the controlled person is successful, the controller will
>experience an error and reorganize.

This presumes that the controller remains accessible and continues to control for the same perceptions as during the interaction. If the controller- controlled person interaction is no longer in effect, the (previously) controlled person might still be able to act (successfully or not) to oppose the delayed effect, without resulting controller error. The "victim" might have been fully insured against theft by deception, and so both victim and thief get what they want.

So, when the conflict arises some time AFTER the controller has finished controlling the controlled person's actions, "real" problematic control of another's actions can have been successful.

Because they realize that it can (often, I claim, based on empirical evidence) be successful, many, many people are concerned about the nature and limits of controlling others' actions and of having their own actions controlled by others -- not only when the outcomes of such interactions are problematic for any of the parties, but also when they aid or make possible ("facilitate") control of some perceptions of the controlled person (and employing reorganization to do this, in some cases). I say that PCT has much to tell those people -- much more than the (TRUE) slogan, "You can't make somebody want what they don't want without causing conflict." We're finally getting beyond the (TRUE UNDER CERTAIN INTERPRETATIONS OF THE TERMS) slogan, "You can't be controlled by anyone else." Hooray!

Best, Greg

P.S. Just got Tom Bourbon (921022.13:45) on not watering down PCT. I agree whole-heartedly! (AT LEAST) 4 CHEERS!!!!

Date: Fri Oct 23, 1992 5:42 am PST
Subject: ON THE BILL AND GREG CONVERSATION

CHUCK TUCKER 921023

RE: THE BILL AND GREG (AND OCCASIONAL OTHER) DIALOGUE
    WTP 921017.0900; 19.0900; 20.1100; 21.0915
    GW 921018; 20; 21 EF 921021.0945 DG 921015 [AWOL]

I am clearly in support of such a dialogue and encourage as many people that can enter into it to do so. I am not willing to enter into the conversation in the same fashion as others on the NET; I choose to gather many posts, read through them, take notes on them, think about them, talk to some of my students and colleagues about what is being stated and then (if I think I can write something useful) write a note. I like the exchanges but I have to slow them down rather than jump into the ongoing conversation. If I don't write something it is not out of a lack of interest (this is the most interesting part of my education) but because I don't believe I have anything to contribute in a substantive way.

From my reading of the recent posts (identified above) I still find more agreement between Bill and Greg than I find differences. To do this I discount what I consider to be an occasional "slip" by both parties in the use of words. When I read these posts over time I find that these "slips" are corrected by Bill or Greg or both although there is not always a co-indication of that accomplishment. As you might guess the "slip" is often made with regard to the word 'CONTROL'. I usually notice it and just "plug" in the definition that is consistent with PCT: CONTROL can only be used about the action of a negative feddback system itself - CONTROL as an action, process, event or any phenomena CANNOT happen BETWEEN systems IF they are negative feedback control systems (although quite disturbing, Rick's solution to the AWOL problem by making the systems into SR systems by killing them makes the point dramatically clear to me - although I would note that after a person has been offically declared dead the body still functions as a negative feedback control system). Thus, what happens between people (excuse spelling errors) is coercing, forcing, physically or cognitively manipulating, influencing, persuading, pleading, bribing, requesting, asking. begging, conning, convincing, rationalizing, indirect manipulating ("rubber banding"), agreeing, committing, taking for granted, assuming, "of coursing," "why notting," promising, pledging, contracting, willing, buying, selling, envisioning, respecting, loving, threating, entraping and the like BUT NOT CONTROLLING.

The best "evidence" (I put words in quotes when I want to write loosely or metaphorically) for the PCT model FOR ME is gleaned from presenting problems to be solved and then see how the various parties propose to solve the problem. Notice the proposals by Bill (921020.1100) and Ed (921021.0945) and you will see PCT in action [I see the exchange between Rick and Dennis as a discussion of some variations but not PCT directly but perhaps I am mistaken on this one - I am awaiting Dennis's answer to RM 921021.1000]. Do you notice how the style and tone of the writing seems to alter when Bill and Ed are writing about some "concrete" problem rather that about "theory or model"? I do. If one wants to answer the question: what can you do with PCT that you can't do with other approaches? the posts about specific problems are the most useful. I also think that they tell me much about the model.

I think there are two "topics" of discussion within the Bill and Greg (etal) conversations that keep coming up and don't get adequately "resolved". One is usually mention as "ideology" while the other is "selling PCT." For me these "topics" usually act as disturbances and seem to "spin" the interactants into conflicting but not very clarifying (for me) action. Here is how I make sense out of these two "topics" (I believe that the interactants will agree with me but don't seem to be able to with each other - let me try and see what happens).

IDEOLOGY

When Bill tells others that they could at least reduce the problems that they have with each other if they would just realize that conflict arises when one person tries to do the impossible with another - tries to CONTROL the other - Greg views this as a "how you ought to behave" prescription that is part of PCT. I DON'T THINK THAT IT IS AND I BELIEVE THAT GREG WILL SEE IT WITH A MOMENT OF RELFECTION ON IT. But Bill believes that conflict is both abhorent and unnecessary in human affairs and he has stated this on a number of occasions but most clearly (to me again) in his Chapter 17 of BCP. Thus, I can see that one might see this as an ideological statement which is part of PCT. I can also find where (to me) he has noted how offended he is by conflict and one might see that as part of the model. But as a "scientific" presentation, I don't think that this "ideology" is part of PCT but I see nothing wrong with stating that one of the practical implications (and a very important one) is the reduction of conflict and a greater appreciation of other people. In this regard, one of my most recent arguments against the S--> NO O --> R formulations is that they do lead to seeing coercion, force and threats of force, bribing, physical manipulation, and killing as appropriate ways (sometimes the ONLY way) to get a person to follow your orders and do what they are told to do for their own good. YES, this is ideological but I believe that I have tons of "evidence" to support my case including statements by the proponents of these formulations (see Dennis's statements about the proponents of "behavioral control theories").

SELLING PCT

I have noted in a previous post that I don't believe that you can convince another to take up and use PCT unless he/she is troubled by and finds problematic his/her current "explanation" of "human behavior". This is what PCT says to me; a person will not reorganize unless seriously disturbed about their current actions. I have found that I am not able to disturb most of my fellow sociologists or social psychologists enough to have them take over PCT; the ones who are interested in it (present company excluded) are those that operate at the "fringe" of the discipline and/or are very troubled by their current theories and are SERIOUS enough to want to overcome their problem. I find that one feature of the CSG associates is that we are serious about understanding human life AND do not compartmentalize this concern. Most of the academics I know see sociology (or any discipline) as something one teaches, reads about, studies, and tells others about BUT does not LIVE. This makes it extremely difficult for them to become disturbed because it is not "important" to them (important in the sense that their model of living systems is just part of their job not their life). So, as Bill has been saying, if it is not important to them in the sense that they are controlling for it, it make very little difference what you do - they can't be disturbed enough to reorganize. My proposal is to forget about them and attempt to incorporate PCT into the lives of younger people and see if they pick up on it and use it. I do find greater interest among my students than my fellow academics.

A PROBLEM

What can I do according to PCT to get all of you to vote for the Democratic ticket on November 3, 1992? What can Bill Clinton and Al Gore do according to PCT to get all of you to vote for the Democratic ticket on November 3, 1992? [Send the answer to the second question to 75300.3155@COMPUSERVE.COM the email line of the Clinton/Gore Campaign]

Best to all,   Chuck

Date: Fri Oct 23, 1992 7:45 am PST
From: Dag Forssell / MCI ID: 474-2580

TO: Hortideas Publishing / MCI ID: 497-2767
Subject: uue and more
Message-Id: 20921023154502/0004742580NA2EM

[From Dag Forssell (921023-direct)]

Thanks for note. May I trouble you to download yourself? This will verify validity of my instructions in the starter.doc I sent you and saves some money in total. I will send you stamps to make up for your dollar charge with next submission to the archives.

In addition to what you can read in the starter.doc, you will find among other info in the csg/Index files:

```
programs/msdos:
uud.exe           23449  DOS enhanced decoder
uue.exe           17062  DOS uuencoder

programs/source:
lfilter.c          7746  Bill Power's Tracking Filter (Dag has no idea)
uud.c             13750  Enhanced decoder, C source
uue.c              4734  Enhanced encoder, C source
```

and in the help document:

```
-----------
get foobar      - get a file called "foobar" from the archives  (ASCII)
uuencode x.arc  - get a binary file called "x.arc" from the archives,
                  in uuencoded form (short form uue also works)
----If you want to retrieve a binary file send the command "uue foobar" or
"uuencode foobar" and the file will be uuencoded before it is sent. You then
need uudecode to extract the binary file from the mail message.---
-----------
```

If I have been clear and accurate, you are now in a position to send a single message to the server and help yourself to uue.c, uue.exe encoded, dem1a.exe encoded and anything else that can be helpful and of interest.

Thank you for your help. I'll appreciate your reading of starter.doc.

Your thread is bearing fruit at last, is it not. Chuck's summaries are very helpful, it seems to me.

Greetings to all your family, Dag

P.S. Received spreadsheet. Will process this weekend and also proof What is man. Expect to visit UCLA libraries and take a look at references Tom posted. Most interesting. A prospective customer told me two days ago about a chapter in a book on organizational planning that talked about control theory. I shall get copy and see if Bourbon's friends are the source. I don't know author or references yet.

Date: Fri Oct 23, 1992 8:24 am PST
Subject: linguistics, strategy

[From: Bruce Nevin (Fri 921023 08:53:33)]   Bill Powers (921022.1200)) --

> I'm really pleased to hear that you're thinking of writing a paper on
> PCT and linguistics.

How could I possibly do otherwise? :-)

> I've been meaning to ask -- in your latest foray into Achumawi (sp?)
> land, did you have any time to check out PCT concepts across cultures?

I'm not fluent enough in either applied PCT or Pit River culture to have identified cultural differences specifically in terms of reference perceptions. I work as a chameleon (or sponge) and have to devote a bit of awareness to noticing what changes in me in that context. The people have been for four or five or more generations in profound conflict about principles and system concepts (at least), so it is not a simple matter of immersion in Pit River culture, and most of my attention was devoted to staying afloat while pursuing my primary aims of finding speakers of the language, establishing good relations with them, maintaining existing relationships, and eliciting more linguistic data. A very busy month!

I am more aware than before how characteristics of the Pit River language persist in the ways in which the people use the English language--their "Indian English" dialect. I have to be circumspect exploring this, as calling attention to differences would immediately be taken as criticism of their "imperfect" or "incorrect" English.

(Tom Bourbon (Thu, 22 Oct 1992 20:57:00 CDT)) --

Thank you for your very well focussed critique of the uselessness of brown-nosing after the manner of Carver-Scheier-Hyland-Lord-Hollenbeck. I'll save your short list exemplifying the problem.

What I am advocating is rather the opposite of "deliberately softening its implications". But clearly there is then a problem with being perceived as "determined to elicit rejection."

If one can get a given paper, say Rick's "hierarchy" paper, rejected from the same journal first for "saying nothing new" and then for "not building bridges," one might be able to appeal to the editor's sense of fairness and decency by pointing out the prohibitive inconsistency.

The former is the current status of Rick's paper I take it, and the latter might be the response if he resubmits with a preface framing the whole in terms of metatheory: here is a negative that no one wants, a theory that describes in multiple places what from another point of view are aspects of a single phenomenon. Historical examples of Copernicus et al. a la Kuhn. Touching on how proponents of earlier views in each case resisted disturbance to their established concepts. Then highlighting the business about describing the hierarchical structure of perception and the hierarchical structure of behavior in two kinds of terms, with negative consequences. Explicitly stating the difficulties that readers with commitments to the standard concepts have made their accomodations to their lack of parsimony, and how they will resist disturbance to those concepts. The piece is thus recast in terms of the drama of scientific (r)evolution. Reviewers and editor must reject it (if they do) on those terms, and not merely in terms of the standard concepts in which their commitments are vested.

It is of course easy for me to speak, not having participated in the great joy and pleasure etc. (as you describe it). And there is no guarantee that it would work. Perhaps even it has been tried, by Bill or others. I don't want to get into the WDYYB (Why Don't You, Yes But) seesaw. In part I am thinking through how I will deal with similar processes when I put my show on the road. I know my friend Tom Ryckman's dissertation on the import of Harris's work for linguistic metatheory is not published because, even though editors in three publishing houses in succession (Bradford/MIT, UChi, and one other) liked it a lot and supported it, they could not get any reviewers to open it up and read it. No review, no publication. Pretty effective.

I don't think it's an explicit conspiracy. Rather, any two people who have in common certain arrangements of higher-level reference perceptions are liable to defend them against disturbance in similar ways. Having such perceptions (concepts, commitments) in common and being competent at defending them is what higher education aims to accomplish, no? Or social membership in general.

Bruce bn@bbn.com

Date: Fri Oct 23, 1992 10:01 am PST
Subject: UUdecoder for MS-DOS

A new uudecoder for DOS machines has just become available, and you can get it from biome as pub/csg/programs/msdos/extrct32.zoo -- I haven't tested it, but those of you who have been looking for decoders may want to give it a try if you get uuencoded binaries from the mail server.
--
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: Fri Oct 23, 1992 10:45 am PST
Subject: Controlling perception of action

[From Bill Powers (921023.1030)]   Greg Williams (921022) --\

>control [your perception of -- I hasten to add!]

There's a problem with insisting too steadfastly on referring to control strictly in terms of perception (even though all we can control is perception). In order to talk about control that involves other people, we have to assume that in controlling our own perceptions we are causing things to happen in the outside world that others can see, feel, etc.. In short, we have to include the physics- model in the discussion. This is necessary even to suppose that other people exist, for any observer.

This means that controlling our perceptions of another's action amounts to a social interaction only if the actions that we perceive have a boss-reality counterpart in the physical actions -- the outputs -- that the other is producing.

I can control my perception of your action simply by moving myself so I see the action from a different point of view -- you're pushing the lawnmower away from me instead of toward me. This alters my perception, but does not change your output. On the other hand, if you're watching me I can control the direction in which you're looking by moving myself, and now the change in my perception (of the direction in which you're looking) DOES have a boss-reality counterpart, a physical change in your direction of looking. Both of these cases can be described by saying that I'm controlling my perception of your action, yet only one case is a social interaction.

In the rubber-band experiment, suppose that the rubber band is kept in a horizontal plane a foot above the surface where the target dot is located, and where the other dot (above which I want the subject's finger to be) is located. I can move my end of the rubber-band to put the subject's finger on the line from my eye to the other dot; doing this will actually control the subject's action as well as my perception of the subject's hand in relation to the other dot. Or I can move my head to make the line from my eye to the other dot pass through the subject' hand. In the second case, I will have controlled my perception of the relationship between the subject's hand and the other dot without having any physical effect on the subject's hand position -- the subject won't have to move the hand at all.

So when we speak of controlling a perception of someone else's action, we are speaking of a social interaction only if that perception corresponds to a physical change in the other's output.

It looks as though in order to talk about social interactions at all, we have to take an epistemological position -- a practical position if not a philosophical one. We have to believe the world-model in which an external reality really does exist, have properties, have characteristics, behave lawfully, and so on. This commits us to the physical model and the neurology model, or whatever models we can have a consensus about.

Best,   Bill P.

Date: Fri Oct 23, 1992 1:39 pm PST
Subject: Time-spanning control

[From Bill Powers (921023.1100)]

Greg Williams (921023) --

>>If the controlled person does not anticipate the future error, the
>>action desired will take place and the controller will be satisfied.
>>When the future arrives, however, the effect on the controlled
>>person (if it occurs as predicted) may cause an error. The
>>controlled person will then act to oppose the effect. If this action
>>by the controlled person is successful, the controller will
>>experience an error and reorganize.

>IF such action is possible. There might be no way to "right the
>wrong" -- no successful (or maybe even partly successful) action possible.

Read one more sentence:

>>If it is unsuccessful, the controlled person will experience an
>>error and reorganize.

>This [the controller experiences an error] presumes that the
>controller remains accessible and continues to control for the same
>perceptions as during the interaction.

If the controller is not around at the finish, where is the control? Don't confuse prediction with control.

The above also assumes that the controllee has not changed circumstances so the future effect is no longer relevant, or has not changed goals or perceptions so the future effect is no longer considered adverse, or has not taken steps to prevent the future effect from happening at all, or has not learned in the interim how to oppose disturbances like those of the future effect, and so on and so on. There are many assumptions involved in the hypothetical scenario when we talk about effects of present actions on future consequences. The future is not a fixed function of the past, at least not as far as any human being knows. We make it up as we go. Controlling over a span of time is very difficult, not to mention very slow (how rapidly can the controller react to counteract failure of control with a year's delay?).

>So, when the conflict arises some time AFTER the controller has
>finished controlling the controlled person's actions, "real"
>problematic control of another's actions can have been successful.

Did I suggest anywhere that it can't be successful on any given occasion? Of course if the controller isn't present to maintain the control, it's hard to see how this could be called "control." The controller can see to it that there's a chance that a future disturbance will occur of a type that the controllee might not be able to handle when it occurs. But this isn't to say that the future effect will actually occur as predicted, or that the controllee won't learn to handle it. The controller can persuade the controllee to sign the deed. But when this proves to have been a mistake by the controllee (a problem arises) the controller can't then get the controllee to sign ANOTHER deed (even if there is no future danger in doing so).

When we introduce the time dimension, we can't arbitrarily cut it off after a single time-spanning experience. Over time, people learn from experiences, mostly from doing things or having things done to them that create error for themselves. I can fool a toddler into thinking I have pulled the end of one finger off, so I can make the toddler laugh (or cry) in this way -- once or twice. But I can't go on doing that indefinitely. In the short term I can use this means to control the toddler's actions, sort of. But in the long term the toddler reorganizes and can no longer be controlled in that way. Repeated experience with situations that lead to error results in reorganization that continues until those situations can no longer cause uncorrectable error, or until death.

>Because they realize that it can (often, I claim, based on empirical
>evidence) be successful, many, many people are concerned about the
>nature and limits of controlling others' actions and of having their
>own actions controlled by others -- not only when the outcomes of
>such interactions are problematic for any of the parties, but also
>when they aid or make possible ("facilitate") control of some
>perceptions of the controlled person (and employing reorganization to
>do this, in some cases).

I think that most people who are concerned in this way are trying to achieve control of something that is at best only loosely and uncertainly controllable. I think that most people who think they HAVE control of others, for good or for evil, are indulging in wishful thinking. I think that people are very, very bad at predicting either good outcomes or bad ones.

There is a great difference between wishing that you could control the future and actually being able to do so. There is a great difference between wishing to facilitate the learning of children and actually having that effect. What seems to one person like aiding or facilitating the control of others may appear to another like meddling and coercing. Remember that judgments about the success of controlling others are usually made by the controllers, not the controllees. And even when controllees agree that the effects existed and were good, they can be mistaken: they may not realize how much the success depended on themselves.

I also think that people quite unnecessarily assume that they are being controlled against their will or without their knowledge, when in fact it's their own assumptions and beliefs that lead them to behave in the way they see as externally controlled. They don't realize that it's their own goals and perceptions that trap them, and that by rethinking their goals and perceptions they could shed the apparent control without difficulty, if it is causing them any problems.
I'd like to see the empirical evidence you have. There may be interpretations of it that lead to different conclusions than those you have decided upon.

Chuck Tucker (921023) --

> what happens
> between people (excuse spelling errors) is coercing, forcing, physically or cognitively manipulating, influencing, persuading, pleading, bribing, requesting, asking. begging, conning, convincing, rationalizing, indirect manipulating ("rubber banding"), agreeing, committing, taking for granted, assuming, "of coursing," "why notting," promising, pledging, contracting, willing, buying, selling, envisioning, respecting, loving, threating, entraping and the like BUT NOT CONTROLLING.

I agree. To carry this debate much further than it has gone until now, I think we have to start looking at all these other kinds of interactions that people often CALL control of others and HOPE is control of others and see what is really going on. Perhaps when Greg presents his empirical evidence we will be able to separate what is control from what is not, and begin exploring other modes of interpersonal behavior.

> My proposal is to forget about [those who don't want to learn] and attempt to incorporate PCT into the lives of younger people and see if they pick up on it and use it.

The only selling we have to do is to make sure that what we present hangs together and fits the facts. Those who are interested in ideas with those qualities will come to PCT. Those who are not are not the sort of people I want to hang out with anyway. Everyone in the CSG and on the net is self-selected. Nobody had to write commercials telling these people how great control theory is, and how much better they would feel if they analyzed two control systems before breakfast every day. There are people who have prepared themselves to grasp the principles of PCT and who have no stake in preserving the sciences of life as they are. If we can get our materials out into public view, those people will see them, understand them, and get involved. That is how people became attracted to my ideas in the first place, how the CSG came into being and grew, how this network conference arose and grew, and how we will progress in the future. We have nothing to sell: only something to describe as well and clearly as we can. If there is something in it, the right people will see it.

> What can I do according to PCT to get all of you to vote for the Democratic ticket on November 3, 1992?

Don't push. This simply creates counterefforts.

Anyway, partisan activism on this publicly-supported net is, I believe, a no-no. Couldn't it cost some not-for-profit institutions who help support bitnet and internet their charters?

I don't support any political party. What I support is anyone who is out to enhance the will of the people and who aims (realistically) to minimize the coercive control of others. If others agree with such PCT-informed concepts, they can certainly make up their own minds as to which candidates come closer to meeting these and other such criteria.

Best to all,   Bill P.

Date: Fri Oct 23, 1992 2:22 pm PST
Subject: Misc replies

[From Rick Marken (921023.0930)]   Bruce Nevin (Thu 921022 10:46:41)

Thanks for the comments on the paper; very helpful, as usual.

I agree with Marcos Rodrigues (921023.1000) that calling a reference signal a command may be obscure, but basically correct; in the model the reference input "commands" a particular level of perception.

>Question: what results does PCT predict that standard theories would
>not, or what results do standard theories predict that PCT does not?

>What are the consequences for standard theories of treating them
>separately, which may be avoided by PCT? Why should they care?

The paper was an attempt to present PCT without being negative about other "standard" theories. I did mention at least one prediction of PCT that would not be predicted by conventional theories; behavioral limits are often perceptual limits, not "output generation" limits. This is a testable prediction of PCT. For example, you could show that a sequence cannot be controlled if it occurs at too fast a rate -- even when their is no "output" limitation on the ability to produce the sequence. So you can show that the finger tap sequence cannot be controlled even if producing the sequence is just a matter of pressing a button. A change in the sequence would require a correction -- another press of the button -- so all the person has to do to control the sequence is press one button with one finger -- something that can be repeated rather rapidly. But I predict that as soon as the sequence becomes too fast to perceive as a sequence, the person will not be able to push the button at the right time to correct it.

I didn't want to go into the "consequences for standard theories" in the paper because that gets us into "negative campaigning" again. People who are able to understand the paper (because they have already read PCT stuff) should see the implications just fine. Recall that I was writing the paper for people who had already passed PCT 101. I think that's the only way to write a PCT paper without saying nasty things about the "other side".

Dick Robertson (921022) --

I seem to have lost parts of your post. It seemed interesting. Could you repost it with narrower margins?

Tom Bourbon (921022.13:45)

I knew Tom when he was just a regular, excellent conventional psychologist -- and I was the firebrand. Readers beware; when you start to understand PCT, this could happen to you.

What can I say Tom: A wonderful, moving post. I just wish I didn't know so well what you are talking about.

For the benefit of voyeurs on this net, let me just explain what I think the

> Carver-Scheier-Hyland-Lord-Hollenbeck.

crowd are missing, and why people like Bandura are so hostile to PCT.

According to PCT, behavior is controlled perceptual variables. In order to control those variables (relative to reference levels determined autonomously, by the organism itself) organisms must act to bring those perceptions to the reference level and counter any disturbances to the perception.

The main goals of research in PCT are to determine 1) what perceptual variables an organism controls and 2) how this control is accomplished. The basic methodology for accomplishing 1) is "the test for the controlled variable" -- a methodology that is quite different than the traditional independent - dependent variable approach of conventional psychology (in fact, PCT shows that the conventional methodology reveals nothing about the nature of the organism; just statistical relationships between disturbances and compensating actions). The approach to accomplishing 2) is modeling. Once we know what an organism might be controlling, we try to build working models that will control the variable in the same way. We consider the model a success when it acts just like the real organism (the "why 99%" criterion).

So PCT goals and methods are quite different than those of conventional psychology. This makes conventional psychologists nervous and hostile because it is a major disturbance to their familiar way of going about their business. The Carver-Scheier-Hyland-Lord-etc crowd have made PCT acceptable by jettisoning the two aspects of PCT discussed above; they don't test for controlled variables and they don't model. They do IV-DV research and use statistics to determine if there was a significant effect of one variable on the other (with typically trashy results -- such as the correlations reported by Tom). If these folks really understood PCT, they would know that they are wasting their time doing what they are doing -- but that would be true of other conventional psychologists too. People don't like to think that they are wasting their time -- so there is an easy solution -- don't understand PCT (in fact, don't even TRY -- hence, the hostility).

Dag Forssell (921022)--

>It occurs to me that the time aspect of HPCT can be portrayed as follows:

```
>                     ------------------------------>>> TIME
>
>Syst Conc ***************************************************************
>Principle                  ***************************
>Program                         *****************
>Sequence                          *************
>Category                             *******
>Relationship                          *****
>Event                                  ***
>Configuration                           *
>Transition                              *
>Sensation                               *
>Intensity                               *
```

I like it. I like it.

Greg Williams (921023)--

I think I get it now. I think it is important for us (PCTers) to be clear about the fact that conflict is only expected when two control systems try to control the same (or very similar) variables. There are cases where we control using other control systems as the means of control (as when we get change from a cashier) -- and there are also cases when we control (successfully) irrelevant side effects of the control actions of others (until it becomes inconvenient to the controllee). So I think the Powers/Williams debate has been very worthwhile; it has certainly helped me -- but I'm trying to write a book about control so it has been REALLY interesting to me. Obviously, you now qualify as a most important reviewer of said book if it ever turns into actual typed copy. Would you be willing to spend some time on it?

Best regards   Rick

Date: Fri Oct 23, 1992 3:26 pm PST
Subject: Re: Modeling reorganization

[Oded Maler (921023)   * [From Bill Powers (921016.0930)]
*
* What kind of experiment could you do to test your proposal that the
*"sorry" loop is lower in the hierarchy than the conscious one?

I don't have for the specific one, but in general you can let someone play with the computer such that he/she has to respond in a certain way to some stimulus (sic.) appearing on the screen (e.g., track it with the cursor/mouse) within some time interval. Then make the stimulus more complicated (a larger bit map) and change the rule of the game such that those bitmaps divide into two categories A and B such that for A you have to respond as before but for B you shouldn't. But you loose much more points if you ignore A than if you "over-react" and respond to B. Now you train the subject and compare how the performance changes with the complexity of the categorization (more complex distinction, e.g., degree of darkness vs. the parity of the number of bits) together with the length of the time interval within which he must respond. Finally you introduce an additional activity in parallel (track completely different objects of the other side of the screen) and see how the intensity of the latter activity influences the over-reaction in the previous one.

But I'm not an experimentalist..

--Oded

Date: Fri Oct 23, 1992 7:21 pm PST
Subject: Re: PCT popularity; Why 99%?

From Gary Cziko 921024.0300 GMT}

I have enjoyed the recent comments by Williams, Marken and Bourbon et al. on "selling PCT."

But isn't there a third way to get PCT a wider audience which doesn't involve either (a) neutering and building shakey bridges to mainstream psychology or (b) presenting it as contrary to everything editors, reviewers, and readers have ever understood about psychology? I am thinking of (c) showing how PCT provides a unified theory of psychology in which the various flavors of psychology now existing (e.g., cognitive, S-R, reinforcement theory, etc.) can be seen as dealing with special, (very) limited subdomains of behavior and cognition. Yes, I am thinking of Marken's "The Blind Men and the Elephant" approach, dressed up to show how PCT can potentially handle the whole darn elephant.

That is why I like Rick's paper so much and want to see this published very badly. But I am puzzled that I seem to be the only CSGnetter who is really excited about the potential of Rick's paper. Even Rick doesn't refer to it in his own comment on PCT (of course, we all know how modest and shy he really is)! Am I missing something wrong with Rick's unification approach? Or are other CSGnetters missing an important way of "selling" PCT?--Gary

Date: Fri Oct 23, 1992 9:29 pm PST
Subject: `negative evidence' in lang. acq.

[Avery Andrews 921024.1519]

A linguistic thought:

It is a standard tenet of current generative grammar that learners don't have access to negative evidence. Here's a story about how this may be wrong.

Suppose that comprehension is effected by a very robust, error-tolerant, context- & content-= sensitive system. If it fails to find a suitable meaning for something that gets said, that generates an error-signal that prompts some reorganization (this is off the main point, & just included to provide a wider context).

But when the comprehension system does find a reasonable meaning for an utterance, something else happens: that meaning is sent to the production system, which in effect returns a list of all the different ways it might express the same meaning (assuming the perspective of the original utterer). If the original utterance doesn't come up on this list, you get a big error signal promoting reorganization of the production system, but you also get an error signal if there turns out to be more than one way of producing the meaning. E.g., if the current production system provides alternatives to what was said.

So if your grammar allows you to say:

  John donated the money to the fund.
or:
  John donated the fund the money.

to mean the same thing, there will be some error-signal generated each time the first is used, which can be eliminated by altering the grammar to exclude the second.

So the non-occurrence of the second (a piece of `negative evidence') becomes accessible to the learning system. A further prediction is that the kind of optionality above will be an inherently unstable feature of languages: if two such forms are used with no discernable difference in meaning, the language-acquisition systems of the speakers will be constantly reorganizing without being able to find an error-free configuration, & presumably at some point one of the forms will other drop out, or they will acquire subtly different meanings (like the emergence of magnetic domains in cooling iron, perhaps).

Avery.Andrews@anu.edu.au

Date: Sat Oct 24, 1992 4:39 am PST
Subject: Epistemology & "Control"

From Greg Williams (921024)

Ed, Bill, Tom, et al.: Bill Williams, Apt. 6, 519 North Gate, Pendleton, OR 97801 (no phone)

Ed, I sent the other address to Mary.

>CHUCK TUCKER 921023

>I am clearly in support of such a dialogue and encourage as many
>people that can enter into it to do so. I am not willing to
>enter into the conversation in the same fashion as others on the
>NET; I choose to gather many posts, read through them, take notes
>on them, think about them, talk to some of my students and
>colleagues about what is being stated and then (if I think I can
>write something useful) write a note.

I appreciate your comments.

>Bill Powers (921023.1030)

>There's a problem with insisting too steadfastly on referring to
>control strictly in terms of perception (even though all we can
>control is perception). In order to talk about control that involves
>other people, we have to assume that in controlling our own
>perceptions we are causing things to happen in the outside world that
>others can see, feel, etc.. In short, we have to include the physics-
>model in the discussion. This is necessary even to suppose that other
>people exist, for any observer.

This "problem" isn't limited to "social interactions": it is the same for control that involves "inanimate objects." Yes, it's all perception -- including the faith that we aren't solipsists (always a possibility, but one I don't take very seriously -- that's my ideology).

>I can control my perception of your action simply by moving myself so
>I see the action from a different point of view -- you're pushing the
>lawnmower away from me instead of toward me. This alters my
>perception, but does not change your output. On the other hand, if
>you're watching me I can control the direction in which you're looking
>by moving myself, and now the change in my perception (of the
>direction in which you're looking) DOES have a boss-reality
>counterpart, a physical change in your direction of looking. Both of
>these cases can be described by saying that I'm controlling my
>perception of your action, yet only one case is a social interaction.

On the basis of the PCT model only -- no ideology -- both examples simply count as instances of a party controlling some of their own perceptions which depend on some of another's actions. The first example corresponds to #1 in my summary (NOT disturbing another's controlled variables, but altering the environment so as to perceive what you want which depends on the other). The second example corresponds to #2 in my summary (disturbing another's controlled variables so as to perceive what you want which depends on the other). The PCT model (which I thought we were trying to get back to) doesn't require any talk of "social interactions," as it doesn't need any talk of "importance."

>It looks as though in order to talk about social interactions at all,
>we have to take an epistemological position -- a practical position if
>not a philosophical one.

Talking about control of one's perceptions which depend on another's actions requires taking no more nor less of an epistemological position than does talking about control of one's perceptions which depend on the "motion" of a "rock."

>Bill Powers (921023.1100)

>If the controller is not around at the finish, where is the control?
>Don't confuse prediction with control.

The controller WAS around at the finish (of control -- of getting what he/she wanted), but is NOT around at the time the other party begins to perceive a problem (and is not controlling, then, for what he/she was controlling before). And, after the finish of control, the controller needn't care whether the other party EVER perceives a problem, since he (the controller) got what he wanted and split. (Note that the same sort of analysis holds for nonproblematic "beneficial" (to the "controlled") future consequences: the "teacher" can be controlling all along for his/her perceptions of certain types of actions by the "student" and be long gone -- maybe even dead -- by the time the student realizes the benefits of the exercises done so long ago.)

>When we introduce the time dimension, we can't arbitrarily cut it off
>after a single time-spanning experience. Over time, people learn from
>experiences, mostly from doing things or having things done to them
>that create error for themselves.

And if you lose your life savings in the first scam you've ever been hooked by, that you will learn to be wary next time is woefully miniscule consolation. That's one reason why many folks are interested in understanding control of one's perceptions which depend on others' actions: to help PREVENT such catastrophes. (Many folks are also interested so they can improve counseling/teaching techniques, too.)

>I'd like to see the empirical evidence you have.

It doesn't matter to me that people who claim to be "controlling others" are wrong about that claim. Even if they are controlling some of their perceptions which depend on others' actions, they certainly can delude themselves about what they are doing. That's beside the point (although it might make an interesting sociological study) that such control is possible (according to PCT) and of interest to many people. However it is interpreted, rightly or wrongly, by whomever, the PCT model says that it is POSSIBLE to control some of one's perceptions which depend on some of others' actions, and the PCT model says that there are LIMITS on this, having to do with the controller having an adequate model of the others' wants and with what the others' wants are. Such control CANNOT be arbitrarily accomplished.

There are ubiquitous examples (look around!) of most everyone controlling for some of their perceptions which depend on others' actions. Analyze your own successful controlling in social interactions over the course of a single day, or sit in on an elementary school class and ask the teacher and students about their successful controlling in classroom situations, or go to an office, or a store. I predict that you will hear some pretty wild hypotheses about who/what is controlling who/what, but you'll also see and hear plenty of evidence of control of individuals' perceptions which depend on others' actions.

CHUCK TUCKER>>what happens
CT>>between people (excuse spelling errors) is coercing, forcing,
CT>>physically or cognitively manipulating, influencing, persuading,
CT>>pleading, bribing, requesting, asking. begging, conning,

CT>>convincing, rationalizing, indirect manipulating ("rubber
CT>>banding"), agreeing, committing, taking for granted, assuming,
CT>>"of coursing," "why notting," promising, pledging, contracting,
CT>>willing, buying, selling, envisioning, respecting, loving,
CT>>threating, entraping and the like BUT NOT CONTROLLING.

>I agree.

So do I. But let's be careful to note what we're agreeing about. Chuck was trying to make the point that a person CANNOT control ANYTHING except (some of) his/her OWN PERCEPTIONS. He was NOT trying to claim that a person can never control (some of) his/her own perceptions which depend on (some) others' actions. Chuck's listed activities CAN, at least sometimes, amount to folk descriptions of instances of such control; what he meant by "BUT NOT CONTROLLING" is that they are NEVER instances of controlling ANYTHING OTHER THAN A CONTROLLER'S OWN PERCEPTIONS (specifically, they are never instances of controlling others).

>To carry this debate much further than it has gone until now,
>I think we have to start looking at all these other kinds of
>interactions that people often CALL control of others and HOPE is
>control of others and see what is really going on.

Let's start looking at instances of control of one's own perceptions which depend on others' actions, REGARDLESS of what ANYONE calls it. I'm sure we can come up with MANY instances of what some people WRONGLY call control (such as "Plato, through his writings, CONTROLS that philosopher's world-view" -- give me a break, Plato is long dead and cannot control ANYTHING!), but I think (ideology, again!) that we will attract more nonPCTers to the PCT model if we focus on what the model has to say about the nature and limits of true control in social situations.

>Anyway, partisan activism on this publicly-supported net is, I
>believe, a no-no. Couldn't it cost some not-for-profit institutions
>who help support bitnet and internet their charters?

Last week on National Public Radio, I heard about some folks having a mock "buy-sell votes" national "election" on Internet, so it needn't be as non- partisan as you think. I'm voting AGAINST Quayle, who once said, "I stand by all the misstatements I've made."

Best, Greg

Date: Sat Oct 24, 1992 11:01 am PST
Subject: SR and "sorry"; control of meaning; social control

[From Bill Powers (921024.0830)]

Oded Maler (921023) --

> ... you can let someone
>play with the computer such that he/she has to respond in a certain
>way to some stimulus (sic.) appearing on the screen (e.g., track it
>with the cursor/mouse) within some time interval.

I'm not sure how you see the idea of stimulus and response entering into tracking a cursor. In our experiments there is always a disturbance, so "responding to the stimulus in a certain way" will not result in tracking. There's no response that corresponds to each target position such that the result will place the cursor at the target. If you consider the cursor as part of the stimulus, then it's neither a dependent nor an independent variable.

It seems to me that all the experimental suggestions you made are the S-R sort of experiment, in which you vary an independent variable and look for a correlated change in the dependent variable. Am I misinterpreting your description?

Perhaps I didn't make my question clear. When I asked what kind of experiment you could do to determine that the "sorry" loop is lower in the hierarchy than the conscious loop, I was only asking what criteria you would use to determine that saying "sorry" was lower in the hierarchy than some conscious kind of behavior, such as evaluating whether saying sorry is appropriate.

>But I'm not an experimentalist.

You must have some sort of actual behavior in mind, and use some way of interpreting what you observe. How else can any theoretical concept have meaning?

Gary Cziko (921024.0300) --

Seems to me you mentioned, and Rick accepted, an idea about making Blind Men a joint paper. Somehow I think that Rick's brilliant idea, coupled with your brilliant strategy, would do something like square the importance of the paper. As well as doing what Bruce Nevin suggested -- trying the same paper with a different slant and leaving the editor the choice between publishing and supporting contradictory judgments of the same idea.

Avery Andrews (921024.1519) --

>Suppose that comprehension is effected by a very robust, error-
>tolerant, context- & content-= sensitive system. If it fails to
>find a suitable meaning for something that gets said, that
>generates an error-signal that prompts some reorganization (this is
>off the main point, & just included to provide a wider context).

So far so good: the higher system controls for meaning by finding A lower-level sentence that produces that meaning. Reorganization takes place when no sentence with a suitable meaning can be found.

>But when the comprehension system does find a reasonable meaning
>for an utterance, something else happens: that meaning is sent to
>the production system, which in effect returns a list of all the
>different ways it might express the same meaning (assuming the
>perspective of the original utterer).

This gets a little awkward from the modeling standpoint -- sending a meaning to the production system implies that the production system knows something about meanings. But it's the higher system that's concerned with meanings, isn't it?

Suppose we say, instead, that there is more than one higher system concerned with meaning. If many higher systems are operating in parallel, each concerned with a different facet of meaning derivable from the same lower-order sentence, you get an effect similar to what you describe: one sentence, many possible meanings. But now the meanings stay at the level concerned with meanings instead of sentences, and the sentence-level systems don't have to deal with meanings.

>If the original utterance doesn't come up on this list, you get a
>big error signal promoting reorganization of the production system,
>but you also get an error signal if there turns out to be more than
>one way of producing the meaning. E.g., if the current production
>system provides alternatives to what was said.

The multiple meaning systems put a slightly different slant on this. Each meaning system wants to perceive a different meaning. The sentence that is found produces meanings satisfying one or more meaning-control systems, but leaves others with errors.

>If the original utterance doesn't come up on this list, you get a
>big error signal promoting reorganization of the production system,
>but you also get an error signal if there turns out to be more than
>one way of producing the meaning. E.g., if the current production
>system provides alternatives to what was said.

If there are meaning-control systems yet unsatisfied, reorganization (or at least a continuation of the search strategy, which need not be random) continues. If the current production system produces UNWANTED meanings, it can CAUSE error in one or more meaning-control systems.

>So if your grammar allows you to say:
>
> John donated the money to the fund.

>or:
> John donated the fund the money.

>to mean the same thing, there will be some error-signal generated
>each time the first is used, which can be eliminated by altering
>the grammar to exclude the second.

Under the meaning-control hypothesis, the second would be eliminated only if it produced a meaning that caused an error in some meaning- control system. For example, there might be a principle that you mention the most important subject first. If the question is "What happened to the money?" the answer would be stated "John donated the money (to the fund)." But if the question is "Who got the money?" the answer would be "John donated the fund the money, not the Treasury." To answer the first question the second way would be to violate the principle by implying that the fund is more relevant to the question than the money. But if the question were "What did John do?" then either answer would supply the wanted meaning, and there would be no reason to eliminate either one.

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

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

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

The implication is that the mature speaker will come to express the same meanings in the same words all of the time, never paraphrasing or varying the wording. I'm not sure you would want to maintain that. Under the multiple-meaning-control version of your hypothesis, the choice of wording will be such as to satisfy all of the meaning- control systems that are involved, if possible. This takes care of your conjecture that different wordings will have subtly different meanings -- by which I presume you mean in addition to the main and obvious meaning of all the different wordings. Reorganization would take place only if no wording were available (through a learned search process) that could make all the alternative meanings fit all of the meaning reference levels at once. I think that a control-system model would work better than an annealing model based on principles of magnetism!

Greg Williams (921024) --

>> In short, we have to include the physics-
>>model in the discussion. This is necessary even to suppose that
>>other people exist, for any observer.

>This "problem" isn't limited to "social interactions": it is the
>same for control that involves "inanimate objects." Yes, it's all
>perception -- including the faith that we aren't solipsists (always
>a possibility, but one I don't take very seriously -- that's my ideology).

Right. The non-solipsism assumption is part of the model.

>On the basis of the PCT model only -- no ideology -- both examples simply
>count as instances of a party controlling some of their own perceptions
>which depend on some of another's actions. The first example corresponds
>to #1 in my summary (NOT disturbing another's controlled variables, but
>altering the environment so as to perceive what you want which depends on
>the other). The second example corresponds to #2 in my summary (disturbing
>another's controlled variables so as to perceive what you want which
>depends on the other).

Let's see, I'm not clear now on #1. If you alter your environment to perceive what you want that depends on another's actions, but altering the environment has no effect on that other, this must mean that you're altering something that only you have a reference level for and that does NOT result in any change in the other's outputs. Would a valid example be looking at a distant person through binoculars? I'm trying to make sure that I understand your use of "not disturbing" -- whether you mean not applying a disturbance at all, or applying a disturbance but within the other's capacity to resist, or only refraining from applying so large a disturbance that some controlled variable is materially altered.

>The PCT model (which I thought we were trying to get back to) doesn't require >any talk of "social interactions," as it doesn't need any talk of "importance."

I think we need the term "interaction" to mean the situation where two systems simultaneously disturb each other's controlled variables in the process of controlling their own. I would defend the use of "social" interaction as specifying that the interaction is taking place between equal organisms (with significant loop gains) rather than between an organism and a rock (which has no loop gain). The nature of these two interactions is quite different.

It's true that "interaction" is not a term from PCT. But when we look at relationships between organisms, the rules of PCT no longer apply: there is no superordinate system controlling for any form in the interaction. We can still model social systems, but the approach has to be more that of the System Dynamics people, with organisms and their PCT properties being the building blocks along with non-living aspects of the environment and their properties.

>Talking about control of one's perceptions which depend on another's actions
>requires taking no more nor less of an epistemological position than does
>talking about control of one's perceptions which depend on the "motion" of a
>"rock."

True. In both cases we have to assume an environment in which things happen that are not necessarily represented -- realized -- in perception. It seems necessary to be able to distinguish between what a person perceives to be happening and what is "actually" happening. Pinning down the meaning of "actually" is pretty difficult, unless we confine ourselves to the physical-world model where disagreements are minimized.

>>If the controller is not around at the finish, where is the
>>control? Don't confuse prediction with control.

>The controller WAS around at the finish (of control -- of getting
>what he/she wanted), but is NOT around at the time the other party
>begins to perceive a problem (and is not controlling, then, for what
>he/she was controlling before). And, after the finish of control, the
>controller needn't care whether the other party EVER perceives a
>problem, since he (the controller) got what he wanted and split.

OK, this eliminates the idea of disturbing a present perception in order to control for a future effect on the person. I'm just as happy to let it go, because if any such control were possible, it would be extremely loose, and might take several lifetimes before an error could be corrected.

>(Note that the same sort of analysis holds for nonproblematic
>"beneficial" (to the "controlled") future consequences: the
>"teacher" can be controlling all along for his/her perceptions of
>certain types of actions by the "student" and be long gone -- maybe
>even dead -- by the time the student realizes the benefits of the
>exercises done so long ago.)

Predicting that teaching will have "beneficial" results is not a control process, except in the imagination at the time the prediction is made. Your note explains the motivation of the teacher, but says nothing about whether beneficial effects can be systematically produced in this way. It's pretty hard to produce a systematic result if you can't detect errors and vary your actions to correct them.

I think you've brought out a major deficiency in the commonsense approach to teaching. You know -- "You should learn algebra now, because some day you'll be thankful that I made you learn it." This approach doesn't work very well, if at all -- think of how many people learn algebra this way, and how many of those have any use for it in later life or even remember it. This suggests that perhaps we should focus on teaching algebra in a way that has observable benefits NOW, that helps a student solve a problem that the student wants to solve. Or else we should postpone teaching it until such a problem arises. Many people will never encounter such a problem. They will be no worse off for not being taught algebra than if they had learned algebra and then forgotten it. And they might have spent their early years on learning something nearer to their lines of interest. It occurs to me also that if we didn't force people to learn algebra at a time when they gained nothing relevant out of doing so, it might not be so hard to learn when the need for it did arise. Teaching people what they don't want to know is a great way to turn them off to any subject.

Most psychologists were forced to pass algebra courses. Perhaps if they hadn't been forced to learn algebra, they wouldn't spend the rest of their lives avoiding anything that requires it, like modeling behavior.

>>When we introduce the time dimension, we can't arbitrarily cut it
>>off after a single time-spanning experience. Over time, people
>>learn from experiences, mostly from doing things or having things
>>done to them that create error for themselves.

>And if you lose your life savings in the first scam you've ever been
>hooked by, that you will learn to be wary next time is woefully miniscule
>consolation. That's one reason why many folks are interested in
>understanding control of one's perceptions which depend on others' actions:
>to help PREVENT such catastrophes. (Many folks are also interested so they
>can improve counseling/teaching techniques, too.)

If you're looking for a way to engineer society so that man-made catastrophes will never occur, I think you will look in vain. Maybe that's what people want, but they won't get it. Such catastrophes are caused by people trying to get what they want through controlling what they perceive. You can't prevent people from trying to get what they want. Laws and punishment don't do it; persuasion doesn't do it; teaching doesn't do it; counselling doesn't do it. People will continue to try to get what they want by any means that works, until they are dead, even if you have them locked up in a prison cell.

>Let's start looking at instances of control of one's own perceptions
>which depend on others' actions, REGARDLESS of what ANYONE calls it.

Right.

I have nothing against trying to cope with manmade catastrophes, or any other kind, both before and after they occur. To do otherwise would be like giving up all attempts to control what matters to us, and I don't think that would happen even if I didn't agree that it's not a sensible thing to do. Coping with catastrophes and recovering from them demonstrate the basic human capacity to control and reorganize.

But how well we do this depends on how well we understand what's going on. The first thing we have to understand is that people will continue to try to get what they want by any means that works, until they are dead. That's just a fact of nature. It is not changed by anything that either other people or the environment can do.

Another fact of nature is that people will change whatever they must change in order to get what they value more or at a higher level. They will change their actions and their goals in order to continue to achieve more important or higher goals. They will never act against their own higher goals -- that is impossible. The nearest they can come to that is to adopt conflicting goals, and then they can't act at all.

If we want to change the way in which another person acts, without getting into conflict with the other, we can only arrange ourselves and the world we perceive so that the other has a different way of getting what the other values most, a way that suits us better. This means that we must also change what matters less to us in order to continue to control what matters more. This process requires us to know ourselves well, and to help the other to the same degree of self- knowledge so that the other can tell us what matters more. It requires mutual consent to change, a mutual sense of advantage in maintaining the interaction.

The only human hierarchy of perception and control that we can understand in real time from top to bottom is our own. The Test for the Controlled Variable is useful for proving control, for demonstrating principles, for helping us to choose the control explanation over other explanations of behavior. It can help us see, here and there, roughly what others are controlling for, or what we would be controlling for if we were behaving the same way. But it can never give us a view of another person's organization comparable in detail or accuracy to the view that the other person can have. We are simply in the wrong place to understand another person fully: we are outside, looking at superficial indications; the other person is inside and knows exactly what perceptions are involved.

This is why I maintain that changing society in any substantial and lasting way requires changing the understanding that people have of themselves, and therefore of others. I think that HPCT provides a useful understanding. I don't think that HPCT is going to work through developing tools and techniques that experts can use on society or its members to make them behave more harmoniously or to prevent them from causing more catastrophes. I think that HPCT provides an understanding which will have its effects on society primarily through being understood by the members of that society. Out of that understanding a new kind of society will emerge (as we have seen in our own little one).

Best to all,   Bill P.

Date: Sat Oct 24, 1992 6:19 pm PST
Subject: Re: `negative evidence' in lang. acq.

[from Gary Cziko 921025.0145 GMT]   Avery Andrews 921024.1519 says:

>So if your grammar allows you to say:
>
> John donated the money to the fund.
>or:
> John donated the fund the money.
>
>to mean the same thing, there will be some error-signal generated each
>time the first is used, which can be eliminated by altering the grammar
>to exclude the second.

I realize that this "uniqueness principle" has also been proposed by others looking at language acquisition, but I can't remember the explanation for linguistic forms which seem to be quite synonomous and yet continue to survive as different ways of saying the same thing. While the second of the two examples you give above is not acceptable to my way of speaking English, the both are fine if you use verbs such as "give" or "send." Doesn't your analysis require you then to argue that these have different meanings in the prepositional and dative forms?--Gary

P.S. I received your paper and have skimmed through it but don't any comments yet (until I spend more time with it) other than "thank you."

Gary A. Cziko

Date: Sat Oct 24, 1992 6:19 pm PST
Subject: Objections to PCT

[from Gary Cziko 921025.0200 GMT]

To Bill Powers, Greg Williams & Tom Bourbon and other interested parties:

When I try to discuss PCT with "mainstream" psychologists, two objections often come up: (a) feedback is too slow for many behaviors; (b) deafferentated animals can behave with no sensory input.

I know that both of these subjects have been discussed on CSGnet in the past, but I wonder if there exists published or unpublished papers which address these isssues with more rigor. If these don't exist, perhaps they should.

In the meantime, I remember Bill having made comments about the "speed" objection (and perhaps the deafferentation objection too) and Tom having made comments about the deafferentation studies and I would appreciate if they could summarize their arguments. From Greg I would expect some arguments that the lowest levels may not be fast enough for feedback control with the control effectuated by higher levels sending output commands and perceiving the results of the commands (although I don't want to start Bill vs. Greg on yet another topic before they've settled the current one).

If these objections against PCT are so common, then perhaps PCTers should have objections ready against the objections. --Gary

Date: Sun Oct 25, 1992 7:31 am PST
Subject: Objections to PCT

[From Bill Powers (921025.0800)] Gary Cziko (921025.0145) --

>When I try to discuss PCT with "mainstream" psychologists, two objections
>often come up: (a) feedback is too slow for many behaviors;
>(b) deafferentated animals can behave with no sensory input.

The first objection is a myth tracing back to cybernetics; the second is a straw man argument.

First objection:

If you think of feedback as something that follows after the end of a behavior, then of course feedback is too slow. When you realize that feedback actually starts at the instant that action begins and continues throughout the action, however, that lag disappears.

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

What laymen don't realize is that even WITH lags, a control system can be stabilized by the use of proper temporal filtering in the loop. William Ashby, who did much to start these myths, was a psychiatrist; he didn't know anything about stabilizing control systems. In fact, a properly stabilized control system with a lag can reach equilibrium in one reaction-time, given good filtering. See my Psych Rev article "spadework" where I lay out the requirements for achieving this by using a "slowing factor."

Another related myth is that a stimulus-response system must be faster than a control system. A properly-filtered control system is always at least as fast as, and usually much faster than, a straight-through SR system using the same components but without feedback. The reason is not hard to understand.

If you want a response proportional to a stimulus without feedback, you have to adjust the gain so that in the steady state, the output is of the required magnitude in relation to the input. The device responds by producing an output that rises exponentially to a final value given a step input (all real devices have at least this kind of lag). The maximum possible speed of response without feedback is thus set by the inherent slowness of the physical device.

If we use negative feedback from the output to the input, we can arrange the feedback ratio so that in the final steady state, when the feedback cancels the input, the output is again of the required magnitude. Now, however, we can greatly increase the amplification in the device itself, which is in the forward part of the loop. This does not speed the device up; it still takes as long as before to reach, say, 90% of the final output value. All it does is make that final output value much larger.

Thus when a step input occurs, the initial response of the device is such as to approach an output value that is many times as great as the desired amount of response. The output begins to rise very much faster than in the case without feedback. If there were no feedback, the result would be an enormous overshoot of the desired output value. But as the output increases, so does the negative feedback. When the final state is reached, the negative feedback is cancelling most of the input, and the output becomes exactly the desired amount with no overshoot at all. But the time taken to reach the final state is only a fraction of what it would be without the feedback.

This is exactly what H. S. Black discovered in 1929. He found that vacuum tube amplifiers with a certain inherent gain and bandwidth could be used to achieve not only far more stable gain but a much wider bandwidth, through the use of negative feedback. This knowledge never got into cybernetics, and thus never got into psychology, and thus failed to inform the mythmakers in these fields. Control engineers didn't read the psychological literature, so they never set the record straight.

The fact is that a design with negative feedback is almost always faster in response than a design without negative feedback.

Deafferentation:

This objection is a straw man. There has been a great deal of grisly and clumsy research aimed at disproving a claim that nobody ever made: that without feedback, there can be no behavior. Taub and Bizzi and others, being ignorant of control theory, misinterpreted what control theorists have to say about feedback, and set out to disprove their own misinterpretation. A simple phone call to the right people could have saved decades of effort and a whole lot of misery.

Consider the following control diagram.

```
              | ref sig
              |
     ------> comp ------->
    |                     |
percept                 error
    |                     |
  sensor               effector
    |                     |
     <-- ext feedback <--action
```

WITH feedback the perception is made to match the reference signal. As a result, when the reference signal changes the output changes so as to make the perception change as required for a match. We see the output as "behavior."

Now deafferent the diagram: remove the perceptual signal:

```
              | ref sig
              |
             comp ------->
                          |
                        error
                          |
  sensor               effector
    |                     |
     <-- ext feedback <--action
```

Will there be behavior when the reference signal changes? Of course there will be; there is still an intact path from the reference signal, through the comparator, to the effector. Loss of the negative feedback will result initially in greatly exaggerated outputs and wild instability; this has been observed by everyone who has studied the effects of lesions and injuries on afferent paths. Loss of feedback doesn't give you NO behavior. It gives you MORE behavior. Control theory actually predicts exactly the kind of thing that is seen when the feedback path is interrupted.

There are, however, other feedback paths to higher centers; visual, tactile, and so on. These are not removed by deafferentation. Also, deafferentation is commonly done by cutting the dorsal roots of the spinal cord; this leaves the "auxiliary" pathways in the ventral roots intact.

In any event, deafferented animals are given a post-operative recovery period of, if I remember right, about 16 days before they are tested. During this time, the higher systems learn to control their perceptions using the above un-fed-back output system, and lower their own loop gain so that the reference signals going to the deafferented system have a smaller range of change. This eliminates the gross instability that resulted from the loss of kinesthetic feedback and provides some semblance of normal behavior. What we see is behavior that is basically open-loop with respect to kinesthetic control, but closed-loop with respect to visual control or tactile control. The kinesthetic control system can no longer resist mechanical disturbances, beyond the amount of resistance created by the elasticity of muscles. A load deflects the limb and the deflection is not corrected. Dynamic stability is poor (in fact, in one experiment by Bizzi that I saw, the animal's forearm was strapped to a pivoted board which had frictional contact with a table underneath it: I suspect that this was required in order to keep the arm from oscillating and overshooting. The authors did not explain why this was necessary in order to demonstrate "unchanged behavior").

The whole deafferentation fiasco was motivated not by a desire to understand how behavior works, but in order to defend the conventional view against the threat posed by control theory, as the researchers understood that threat. If these researchers had bothered to study control theory first, to see how it would explain the behaviors they were studying, they would have realized that control theory fits the observations very well indeed, whereas to make conventional theory fit them a great deal of cheating is needed.

Best,   Bill P.

Date: Sun Oct 25, 1992 2:34 pm PST
Subject: negative evidence

[Avery Andrews 921026.0928]
 (Gary Cziko 921025.0145)

>I realize that this "uniqueness principle" has also been proposed by others
>looking at language acquisition, but I can't remember the explanation for
>linguistic forms which seem to be quite synonomous and yet continue to
>survive as different ways of saying the same thing. While the second of

The usual explanation is that the forms aren't really in free variation: there are circumstances that call clearly for one, others that call clearly for the other, and they only look like free variants in the decontextualized `can you say this' interrogation setting.  But I think it is pretty clear that in certain settings, there is often free choice between forms.  So Labov finds that in various `r-less' urban dialects, people put the `r's in in formal settings, leave them out in informal ones, and fluctuate in intermediate cases (or so I think I remember it).

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

Avery.Andrews@anu.edu.au

Date: Sun Oct 25, 1992 2:39 pm PST
Subject: language acquisition

[Avery.Andrews 921026.0936]   (Bill Powers (921024.0830))

>>Suppose that comprehension is effected by a very robust, error-
>>tolerant, context- & content-= sensitive system. If it fails to
>>find a suitable meaning for something that gets said, that
>>generates an error-signal that prompts some reorganization (this is
>off the main point, & just included to provide a wider context).
>
>So far so good: the higher system controls for meaning by finding A
>lower-level sentence that produces that meaning. Reorganization takes
>place when no sentence with a suitable meaning can be found.

I don't think we're talking about quite the same thing here. The higher level system here is supposed to work by controlling for a perception of finding (appropriate) meanings in what people are saying - no sentences are getting produced here.

>>But when the comprehension system does find a reasonable meaning
>>for an utterance, something else happens: that meaning is sent to
>>the production system, which in effect returns a list of all the
>>different ways it might express the same meaning (assuming the
>>perspective of the original utterer).
>
> This gets a little awkward from the modeling standpoint -- sending a
>meaning to the production system implies that the production system
>knows something about meanings. But it's the higher system that's
>concerned with meanings, isn't it?

I'm thinking of the production system as a sort of transducer that turns `meanings', some sort of internal structure, into an overt performance.

Perhaps we should pause to get this ironed out before proceeding.

Avery.Andrews@anu.edu.au

Date: Sun Oct 25, 1992 7:31 pm PST
Subject: Meanings and sentences

[From Bill Powers (921025.1900)]   Avery Andrews (921026.0936) --

>>So far so good: the higher system controls for meaning by finding
>>A lower-level sentence that produces that meaning. Reorganization
>>takes place when no sentence with a suitable meaning can be found.

>I don't think we're talking about quite the same thing here. The
>higher level system here is supposed to work by controlling for a
>perception of finding (appropriate) meanings in what people are
>saying - no sentences are getting produced here.

"Finding" a meaning means perceiving one, doesn't it? And perceived meanings are found from perceived sentences, aren't they?

I'm really fumbling my way through this, because there are so many relationships that I don't understand. Here's a rough picture of what I'm visualizing -- see if it fits your ideas, or could be made to do so.

```
                             desired
                             meaning (nonverbal perception)
                               |
          perceived meaning-->COMP-------error --->
               |                                   |
               |                                   |
         input function                       output function
           |  |  |  |                           | | |  |
           sentences                         sentence-variations
              |                                     |
               <---Imagination or real words <-----
```

Going up the left side, you perceive sentences, which then give rise through some sort of input function to perceived meanings. The perceived meanings (which are not in words, but in terms of perceptions of the things that the words mean) are compared directly with the desired meanings. The error acts through an output function to vary the construction of sentences. These are the sentences that are perceived, closing the loop. Of course when you close the loop by actually uttering the sentences and hearing them, other people hear them too, which takes care of producing utterances.

I would assume that our equipment for perceiving sentences and meanings become organized at least in important ways by learning to create them.

There are clearly more questions than answers in this arrangement. I'm just trying to get across the general PCT concept that we control perceptions, inputs not outputs, by VARYING actions. An error in meaning results in a correction of sentence structure such that the new sentence is perceived and gives rise to a new perceived meaning, which one hopes is closer to the desired meaning. Instead of seeing the process of converting meanings to sentences as an output process, we see it as a process of controlling inputs, where sentence structure and meaning are both perceptions, not outputs.

To make a viable model out of this, it's necessary to specify the kinds of process that have to go on in the input and output functions. The input function must involve conversion from symbol structures into nonverbal experiences, so memory and association would be needed. The meaning-errors would have to be specified in terms of the dimensions of the meaning perceptions -- this is really a schematic of many control systems acting in parallel, each concerned with a different dimension of meaning. Then we have the problem of how to convert a difference between an intended meaning and a perceived meaning (in one of the dimensions) into an appropriate kind of change in the sentence -- that is, a change that will bring the perceived meaning of the new sentence closer to the reference-meaning.

I think the same questions are really implicit in a command-driven model without the feedback loop -- except that the conversions have to go in the other direction. Something has to convert an intended meaning into a sentence that has that meaning. Once you see the feedback way of doing this, the other way (generative) shows its holes, doesn't it? And without the feedback loop, generation of meaningful sentences has to be treated separately from perception of the meanings in given sentences. The feedback model uses the same perceptual system in either case.

>I'm thinking of the production system as a sort of transducer that
>turns 'meanings', some sort of internal structure, into an overt
>performance.

That's always seemed to me to be the weakness in the top-down kind of model. Given a meaning to express, there are very large numbers of sentences that will convey it. What kind of transducer can convert from a single signal into "the right" output among many outputs -- without any feedback to check if that was indeed the right output?

This is the same weakness that exists in the general traditional top- down concept of the brain, where high-level systems issue general commands, which then -- somehow -- get turned into just the muscle tensions that produce just the consequences that could be viewed as appropriate executions of those commands. How do the lower-level systems know that what they are doing, in the current environment, will turn out to be a valid instance of the commanded result? What tells them that moving to the right will execute the command "mail the letter?" The command, being general, can't contain that information. Where does the information come from? This "final common pathway" concept is just not bloody likely.

The control-system model completely takes care of this problem. It removes all the ambiguities in a top-down command-driven model. What is varied first is the sentence; that (going upward in the perceptual systems) varies the meaning; that creates an error that motivates a change (going downward now) in the perceived sentence. This process continues until the perceived meaning matches the desired meaning. Any sentence that does in fact produce the desired meaning is "correct". The only reason it might be changed more is that there is another meaning-control system monitoring for some other aspect of meaning evoked by the same sentence. Then the process of varying the sentence will continue until both meanings are perceived.

When you're just passively listening, the sentences evoke the meanings in the upgoing paths, without the errors (if any) leading to corrective action. Or maybe the kind of corrective action changes -- you say "Do you really mean I should turn right, into the river?"

There are obviously some problems in working out a plausible instance of such an arrangement. How does an error get turned into the right kind of change in a sentence? It doesn't have to produce the correct sentence in one jump, but it has to alter the sentence in a direction that lessens the error. I can think of some sorts of errors, like wrong person, wrong order of events, wrong outcome, incorrect adjective or adverb, mistaken referent, and so on. These would be the dimensions in which the meaning can change, and I would assume that we learn how to vary meanings by altering sentence structure and word choice; that would convert directly into making errors get smaller by varying sentence structure.

I can see where exploring this model would require some novel experiments with language and meaning. It would be informative to see how people correct sentences that convey wrong meanings to them (even if the sentences are grammatically constructed). For a particular kind of meaning error, what sorts of changes in sentences do people make in order to correct them? This would begin to show the dimensions of meaning-error signals, and the dimensions along which sentences can be varied to alter particular kinds of meanings.

This takes a completely new approach to understanding how language and meaning hook together, I suppose. I'm not the person to do it. But whoever does do it is clearly going to have to break with old habits of analysis. The first habit to go would have to be the idea that meanings produce sentences produce utterances. In PCT the chain goes in the other direction... it's all perception.

>Perhaps we should pause to get this ironed out before proceeding.

That may be somewhat optimistic, but we can try.

P.S. Given three nouns A, B, and C, you can create the sentences

The A B'd the C, and
The B A'd the C.

The captain hogged the ice-cream, and
The hog captained the ice-cream.

The only way to judge that one is allowable in ordinary discourse and the other is not is through the meanings of the words. That can't have anything to do with language itself, because the fact that hogs don't captain ice-cream is an accident of experience; given other experiences, we might see immediate sense in the second sentence but not the first, depending on what perceptions "hog" and "captain" are attached to, and what we can imagine hogs and captains doing. There's something about the way words are presented that connects to the way meanings are evoked in juxtaposition and sequence -- that's the input function above. But once the meanings are evoked, the rules that apply aren't those of language, but of perceived reality in general. We don't judge that hogs can't captain things because of language, but because of our experiences with hogs and captains.

That must sound like utter nonsense.

Best,   Uncle Bill P.

Date: Mon Oct 26, 1992 3:04 am PST
Subject: Finis (unless...)

From Greg Williams (921026)   >Gary Cziko 921025.0200 GMT

>... I don't want to start Bill vs. Greg on yet another topic before
>they've settled the current one...

The main issues in the "control which depends on others' actions" discussion have been settled to my satisfaction. If some of the folks who, publicly or privately, expressed interest in that discussion are willing to play active roles in further discussion, I might be persuaded to post more on this topic.

Best,   Greg

Date: Mon Oct 26, 1992 4:03 am PST
Subject: Re: Meanings and sentences

Forgive me because I don't have the original sentences at my disposal right now, but I remember that they were of the type that contain direct and indirect objects.

It is a mistake to judge synonymity at the sentence level. As an obvious example, consider:

Give the book to me.
Give it to me.

Clearly the correctness of the second sentence normally depends on something to which the 'it' refers having been referred to previously or being evident from nonlinguistic context. So, two sentences which may in fact be synonymous each have a different meaning(?) in the larger context.

This is also often true for what many would call stylistic options. So,

Give the book to me.
Give me the book.

may depend on what is being focused on and what is old information. The former, intonation being "normal" would mean that you should give the book to me, not to someone else. The latter would mean that what you should give me is the book, not someone else. (The interplay of intonation, context, and what has been said or done before makes this interpretation in fact more complex than I've just stated, but I think the idea is clear.)

Therefore, two sentences which seem to be synonymous may be so at one level, but are not in actual use.

Most speakers/writers are of course not immediately aware (or perhaps not ever aware) of the distinction, but if you study texts (spoken in particular) it becomes fairly evident that there are distinctions. It is only when one looks at prescriptivist grammar and style that it seems that there is synonymy because such systems are not normally designed to consider context and usage.

Best, Eileen Prince Northeastern University

Date: Mon Oct 26, 1992 4:22 am PST
Subject: Simulating Societies '93

Call for papers and participation

Simulating Societies '93 24-26 July 1993

Approaches to simulating social phenomena and social processes

Although the value of simulating complex phenomena in order to come to a better understanding of their nature is well recognised, it is still rare for simulation to be used to understand social processes. This symposium is intended to present original research, review current ideas, compare alternative approaches and suggest directions for future work on the simulation of social processes. It follows the first symposium held in April 1992 at the University of Surrey, UK.

It is expected that about a dozen papers will be presented to the symposium and that revised versions will be published as a book. We are now seeking proposals for papers and for participation. Contributions from a range of disciplines including sociology, anthropology, archaeology, ethology, artificial intelligence, and artificial life are very welcome.

Papers on the following and related topics are invited:

* Discussions of approaches to the simulation of social processes such as those based on distributed artificial intelligence, genetic algorithms and neural networks, non-linear systems, general purpose stochastic simulation systems etc.
* Accounts of specific simulations of processes and phenomena, at macro or micro level.
* Critical reviews of existing work that has involved the simulation of social processes.
* Reviews of simulation work in archeology, economics, psychology, geography, demography, etc. with lessons for the simulation of social processes.

* Arguments for or against simulation as an approach to understanding complex social processes.
* Simulations of human, animal and 'possible' societies.

'Social process' may be interpreted widely to include, for example, the rise and fall of nation states, the behaviour of households, the evolution of animal societies, and social interaction.

Registration, accommodation and subsistence expenses during the meeting will be met by the sponsors. Partic ipants will need to find their own travel expenses. Proposals for papers are initially invited in the form of an abstract of no more than 300 words. Abstracts should be sent, along with a brief statement of research interests, to the address below by 15th March 1993. Authors of those selected will be invited to submit full papers by 1st June 1993. Those interested in participat ing, but not wishing to present a paper, should send a letter indicating the contribution they could make to the symposium, also by 15th March 1993.

The organisers of the Symposium are Cristiano Castelfranchi (IP-CNR and University of Siena, Italy), Jim Doran (University of Essex, UK), Nigel Gilbert (University of Surrey, UK) and Domenico Parisi (IP- CNR, Roma, Italy).

The symposium is sponsored by the University of Siena (Corso di laurea in Scienze della Comunicazione), the Consiglio Nazionale delle Ricerche (Istituto di Psicologia, Roma) and the University of Surrey.

The meeting will be held at Certosa di Pontignano near Siena, Italy, a conference centre on the site of a 1400AD monastery.

Proposals should be sent to: Prof Nigel Gilbert, Department of Sociology, University of Surrey, Guildford GU2 5XH, United Kingdom Tel: +44 (0)483 509173 Fax: +44 (0)483 306290 Email: gng@soc.surrey.ac.uk

Date: Mon Oct 26, 1992 5:58 am PST
Subject: Statistics referred to in Tom's post of 921022.13:45

A TABLE SHOWING THE RELATIONSHIP BETWEEN SEVERAL DESCRIPTIVE STATISTICS*

r	r2	k2	k	E
1.00	1.00	.00	.00	100 %
.9995	.999	.001	.032	97 %
.9987	.997	.003	.054	95 %
.995	.99	.01	.099	90 %
.954	.91	.09	.299	70 %
.90	.81	.19	.435	56 %
.87	.756	.244	.493	51 %
.865	.748	.252	.50	50 %
.80	.64	.36	.60	40 %
.75	.56	.44	.66	34 %
.71	.50	.50	.70	30 %
.65	.42	.58	.76	24 %
.60	.36	.64	.80	20 %
.55	.30	.70	.83	17 %
.50	.25	.75	.87	13 %
.45	.20	.80	.89	11 %
.40	.16	.84	.92	8 %
.35	.12	.88	.94	6 %
.30	.09	.91	.95	5 %
.25	.06	.94	.97	3 %
.20	.04	.96	.98	2 %
.15	.02	.98	.99	1 %
.10	.01	.99	.995	0 %
.00	.00	1.00	1.00	0 %

DEFINITION AND INTERPRETATION OF THESE STATISTICS**

All of these measures describe two variables (X, Y) within a particular sample:

r is a correlation (or coefficient of correlation) which describes the linear association of one variable with another. It can also be characterized as "... a relative measure of the degree of association between two series " of values for two variables. It varies between 1 (perfect positive correlation) to -1 (perfect negative correlation). The closer this measure is to a perfect correlation the more confidence one has in "predicting" the values of one variable from another variable.

r2 is a measure of "explained" variance (or coefficient of determination) which describes "shared" variation or the amount of variance that one variable is "explained" by the other variable or the proportion of the sum of y2 that is dependent on the regression of Y on X. The larger the numerical value of this measure the more confidence one has in "predicting" the values of one variable from another.

k2 is a measure of "unexplained" variance (or coefficient of nondetermination) which describes "unshared" variation or the amount of variance that one variable is NOT "explained" by the other variable or the proportion of the sum of y2 that is independent of the regression of Y on X. The smaller the numerical value of this measure the more confidence that one has in "predicting" the values of one variable from another.

k is a measure (called coefficient of alienation) which describes the lack of linear association of one variable with another or the ratio of the standard error of estimate to the standard deviation of the variable. The smaller the numerical value of this measure the more confidence one has in "predicting" the values of one variable from another.

E this measure is computed by (1-k)100 and is called an "index of forecasting efficiency" (Downie and Heath, 1965: 226) and indicates the "improvement" for a prediction by knowing the coefficient of correlation (r) for two variables as contrasted with knowing nothing about the linear association of the two variables. For example, with a coefficient of correlation of .71 one can "predict" the values of one variable from another 30% better (on the average) than one could "predict" those values WITHOUT any knowledge of the relationship between the two variables OR one has decreased the size of the "error of prediction" by 30% (on the average) by knowing that the correlation of the two variables is .71.

REFERENCES

Arkin, Herbert and Raymond R. Colton. 1956. Statistical Methods. College Outline series, Forth Edition, Revised.

Downie, N. M. and R. W. Heath. 1965. Basic Statistical Methods. Second Edition. New York: Harper and Row.

*compiled by Charles W. Tucker with the encouragment and assistance of the Control Systems Group CSG-L @ UIUCVMD

(especially Gary Cziko) and the comments of Jimy Sanders.
Other comments appreciated - N050024 AT UNIVSCVM.BITNET

**It should be noted that these descriptions and interpretations, especially those involving "predictions" are limited to a particular sample; if another sample is not a random sample from the same population then predictions about the other variable ("Y") will be unpredictably worse than the original sample.

Date: Mon Oct 26, 1992 7:12 am PST
Subject: Bandwidth

[from Gary Cziko 921026.1205 GMT] re. Bill Powers (921025.0800)

Bill,

Thanks so much for your response to my speed and deafferentation questions. And thanks for your patience. I know you've said all this stuff (probably many times) before, but it may take a while (and a few repetitions) before it sinks into the brains of us non-engineer-background PCTers. I'm sure many people on the net found it useful as well.

But there is a word which often creeps up in technical discussions (especially between you and Martin Taylor) which I don't yet feel I have a good understanding of--bandwidth. To wit:

>This is exactly what H. S. Black discovered in 1929. He found that
>vacuum tube amplifiers with a certain inherent gain and bandwidth
>could be used to achieve not only far more stable gain but a much
>wider bandwidth, through the use of negative feedback.

I have a pretty good idea of how radio works (at least a better understanding than Fred's in Greg Williams's _Fred & Bill_ story), so for me bandwidth refers to the amount of frequency spectrum a signal takes up--for example broad-brand FM (normal commercial broadcasting of moldy oldies) takes up more bandwidth than narrow-band FM (as used in police and amateur radio communications). But when you guys talk about it, you seem to refer to the amount of information being transmitted. In radio, the amount of information is the same from the perspective of words, but more from the perspective of fidelity (normal FM broadcasting has much higher fidelity than narrow-band).

Perhaps you could apply the bandwidth concept to an ECS and its meaning might become clearer for me.--Gary

P.S. Greg, thinking of Fred and Bill made me think of the fast approaching Christmas/New Year's season. Is this because you posted this story last year around Christmas? Anyway, I like this story so much that I would like it to become a Christmas/New Year's tradition on CSGnet. Perhaps you could even make a few minor changes so that it fits the season better.

Gary A. Cziko

Date: Mon Oct 26, 1992 12:38 pm PST
Subject: Re: Bandwidth

[Martin Taylor 921026 12:30] Back today only, then gone for a week or so.
(Gary Cziko 921026.1205)

Gary,

Bandwidth is an interesting issue to bring up. In a linear system, it is quite easy to deal with, in a simplified form, but even in linear systems there are subtleties that are not always evident. In a non-linear system the concept is almost metaphoric, but useful. I have tended to mix the technical with the metaphoric in my postings, because I know of no easy corresponding concept other than informational channel capacity that works in a non-linear system. I'll try to give a brief explanation that applies to a linear system, and to indicate why it matters.

We start with the idea that things vary over time, and we will consider only a scalar variable x (x can at any moment be described by a single real number between minus infinity and plus infinity). The variation over time of this variable is called its waveform, x(t). A couple of hundred years ago, Fourier proved that x(t) can be described by a unique infinite sum of sinusoidal variables. There are an infinite number of these, because to specify x(t) over infinite time one needs to use all frequencies between zero (a steady level) and upward further than any specified number. Frequency is the number of oscillations of the sinusoid per second, so a high frequency means that rapid variations in x(t) can be described. The set of parameters that specify the magnitude and phase of the sinusoid at each frequency needed to match x(t) is called the Fourier Transform of x(t), and is sometimes written as X(f).

If x(t) has no moments of very rapid variation, X(f) will have magnitude near zero for all frequencies higher than some maximum value Xc (should be a subscript c). The value of Xc is called the cutoff frequency of X(f) ("cutoff" is sometimes used in other ways, so don't let it confuse you).

Certain kinds of waveform for x(t) have interesting and important Fourier transforms. One, in particular, is central in our discussions. Call it w(t) (w = white noise). W(f) is zero above a certain frequency Wc, and has a uniform magnitude for all frequencies below Wc. W(f) is said to have a rectangular spectrum. In addition, if you look at w(t) (that is, you sample w(t)) at regular intervals spaced further apart than 1/2Wc the values you get have zero correlation (at this point, that's a further part of the definition of w(t), not a consequence of the shape of W(f)). One sample provides no information about the distribution of values of the next. If you sample at time intervals closer than that, the successive samples are necessarily correlated, and you get no more information about the waveform w(t) than you would have if you sampled at an interval 1/2Wc. Sampling at 1/2Wc is fast enough to tell you everything there is to know about the waveform w(t), or about any other waveform that uses no frequencies about Wc. 1/2Wc is called "the Nyquist limit" or "the Nyquist rate" after its discoverer.

Most signals are not like w(t), in two ways. Firstly, in real signals there is no fixed value of Wc above which the magnitudes of the frequency components are all exactly zero. Rather, the magnitudes tail off over some range until they become too small to detect. Secondly, the magnitudes below Wc are usually not all the same. There are peaks and valleys in the magnitudes as a function of frequency.

So far, I did not mention phase, but it is important, because it determines the actual shape of the waveform x(t). If the phases (the moment at which the sinuoid crosses zero can be used as an indicator) are in some easily described relation to each other, then even if the spectrum is rectangular, like W(f), there will be correlation among the successive samples of w(t) even if the sampling is slower than the Nyquist rate. The waveform will convey less information than will w(t). Of all signals with a Fourier transform contained entirely within a cutoff frequency Xc, white noise (w(t)) requires the most information to describe it. The amount of that information depends on the precision with which each successive sample must be described.

In dealing with signals of less randomness than w(t), it is conventional to talk about "equivalent rectangular bandwidth", which is the bandwith (i.e. the cutoff frequency) of a white noise signal that would require the same amount of information to describe it. So, when we talk about bandwidth, we are talking about two things: (1) the fastest rate at which the samples of an equivalent white noise could be sampled without forcing correlation among the values of successive samples, and (2) the rate at which the signal can convey information, which depends not only on the sampling rate of the equivalent white noise, but also on the precision with which those samples would have to be specified. In the second sense, we mean "information rate" or "channel capacity" whereas in the first sense we are dealing with time- related phenomena that will show up in specific application. For PCT purposes, the most important is probably the intrinsic transport delay around a feedback loop, which affects the stability of the loop. The wider the bandwidth, the faster the loop.

The concept of bandwidth is often applied to filters. In this case, it means that the filter reduces the magnitude of frequencies outside the "filter band." An equivalent rectangular filter with a cutoff Fc would reduce to zero all the frequency components of

X(f) higher than Fc, leaving a waveform x'(t) that could not carry any more information than a white noise with cutoff Wc = Fc.

When we come to non-linear systems, the Fourier transform approach no longer applies. But the concept of information transfer does apply, and so do the questions of loop delay and stability. So we continue to use the word "bandwidth" even though it is technically wrong to do so. However, it is a good metaphor, and much of the time you won't go too wrong by thinking of the effects of bandwidth in linear systems. Sometimes you will go wildly wrong, but that's life.

Hope this helps.   Martin

(Today, according to the polls, we in Canada commit national Hara Kiri. For what perception of honour does this control, Greg?)

Date: Mon Oct 26, 1992 12:54 pm PST
Subject: Re: PCT popularity; Why 99%?

[From Rick Marken (921026.0930)]   Gary Cziko (921024.0300 GMT)

>I am thinking of (c) showing how PCT provides a unified theory of psychology
>in which the various flavors of psychology now existing (e.g., cognitive,
>S-R, reinforcement theory, etc.) can be seen as dealing with special,
>(very) limited subdomains of behavior and cognition.

I do think this is a great idea. But there is still one little, teensie, weensie problem with this friendly approach; it is hard to keep your audience from noticing that their "flavor" of psychology is being revealed as a misinterpretation. The "Blind men" paper doesn't just show that the cognitive, S-R and reinforcement data can be seen as "special cases" of control; it also shows that the explanations of these data are wrong. For example, cognitive models (like the language models we've started discussing again on the net) are output generation models; the appearance (that cognitive behavior is generated output) is taken at face value and so the explanations have been completely off-base. So when you say that PCT is a way of integrating different approaches to understanding behavior, you are also saying that the conclusions that were based on that approach are wrong. Variables aspects of sentences, for example (like meaning, structure, inflection, etc) are controlled INPUTS, not generated outputs. S-R relationships don't depend on the nature of the organism (S-O-R) but on the nature of the environment. Reinforcement is controlled by (not in control of) action. So, while PCT does explain why we see behavior as we do, you can't get away from the fact that it changes how we explain what we see; there is no getting away from the fact that PCT says that psychologists can take their explanations of behavior and toss them in the waste basket. In terms of the "Blind men" analogy, it means that the fellow who has developed all those terrific models of the elephant that explain its "wallness" can now just throw those models away -- the elephant's "wallness" is just a side effect.

I think you have to take this into consideration, Gary, if you want to try to make the "Blind men" paper palatable. It's true that the feedback analysis integrates observations -- and I can see that that fact can be presented to psychologists in a friendly way. But it also shows (and there is just no getting around this) that current models of behavior (that were based on taking these observations at face value) are just, flat out, downright wrong. I hope you can find a nice way to explain THAT. Of course, we could opt NOT to explain that. But then, I think, you risk the "so what" phenomenon. "So what if PCT let's me see SR,cognitive and reinforcment as special cases of a bigger picture; I'm only interested in reinforcement so I'll just keep doing my operant conditioning studies." I think the paper already gets this "so what" response; because it is only in a little section at the end that I gently suggest that ths PCT point of view requires a whole new approach to explaining SR,cognitive and reinforcement phenomena -- an approach based on recognition that what you are seeing is control of perception. So, I would really be interested in seeing what you have in mind to make PCT an attractive option for conventional psychologists.

Best regards   Rick

Date: Mon Oct 26, 1992 1:57 pm PST
Subject: error of type in social control

[From: Bruce Nevin (Mon 921026 11:26:13)]

I've been thinking about this "social control" issue as a problem of logical typing. I'm not sure how far this can get even as a summary of the discussion, but here goes.

When I control a perception P, the behavioral outputs or actions A by which I effect control are typically not themselves perceived as part of P, and often may not even be perceived at all.

The actions A' of another person commonly are included in perceptions that I control (are in P for me).

Thus, when I control for the other person X doing thus and so (for example, contributing to thus and such consequences in my environment) I seek to manipulate something that X is not actually controlling. The (perceptions of the) "same" behavioral outputs A' are perceptions of different logical type for myself and for X, respectively. (Of course, X's perceptions of the actions include many perceptions available to me only in imagination, e.g. kinesthetic perceptions.)

If my influence on some of X's actions does not reduce the net efficacy of all of X's actions in producing results that X *is* controlling, X does not resist my influence and probably does not even perceive it.

A typical scam attracts X's attention to controlling some perception with high gain and therefore ceasing to attend to other controlled perceptions as closely. In a more subtle variation, one could assist X in controlling for P, in such a way as to offset disturbances to P due to one's manipulation of A. We have all had people offer to help with ulterior motives.

There is an added complication in that for normative, convention- oriented behavior we do include aspects of our behavioral outputs or actions A among our controlled perceptions P. Thereby hangs a long tale that we have touched on in the past and I am sure shall again, having to do with why and how we perceive (recreate?) social norms and orient our own behavioral outputs to them. What is germane here is the possibility of manipulating our own actions according to conventional norms so as to trick another person into perceiving our relative status and relationship wrongly, but in a way advantageous to ourselves.

Ordinarily, the details of the behavioral outputs that we call generically body language are beyond conscious manipulation, not just in humans but in all mammals, and probably in other kinds of critters as well. Bateson suggests that there is an evolutionary basis for this. They are the means by which we make ourselves predictable to one another. If they were consciously manipulated by members of a group, they could not serve this purpose and the group (and its members) would not survive as well.

I think it is important and informative to us that members of a social group willingly make themselves more predictable to one another in ways that we talk about as social norms or conventions. They evidently maintain these pre-established agreements (of little direct consequence in themselves) as a basis for more substantive ad hoc agreement and cooperation. The alternatives, generalized as social isolation, ostracism, banishment, appear to evoke great anxiety, at least among mammals, and this suggests to me intrinsic error.

I think it is important and informative to us that almost universally the behavioral outputs by which we enact social norms are kept out of awareness, deliberately not subject to conscious manipulation. They are controlled perceptions subsumed under system concepts (I think that's the right level) like "being a member of this group." And we are very sensitive to discrepencies and inconsistencies that might betray an outsider trying to "pass" as a member.

(Bill Powers (921024.0830)) --

>If we want to change the way in which another person acts, without
>getting into conflict with the other, we can only arrange ourselves
>and the world we perceive so that the other has a different way of
>getting what the other values most, a way that suits us better. This

>means that we must also change what matters less to us in order to
>continue to control what matters more. This process requires us to
>know ourselves well, and to help the other to the same degree of self-
>knowledge so that the other can tell us what matters more. It requires
>mutual consent to change, a mutual sense of advantage in maintaining
>the interaction.

Hear, hear! But be it noted that many scams depend upon the mark being a "nice guy" and accommodating the perpetrator in just this way. And many garden-variety misunderstandings (perhaps most) depend upon differences in how the participants parse the "same" behavioral outputs as constituting quite different communicative acts (e.g. a reminder as part of responsible teamwork, vs. nagging). Knowing oneself involves not just knowing what matters more or less than what else, but also knowing what constitutes an instance of a perception that matters. This is not a given of the environment, it is a matter of ongoing negotiation of agreements. It is no wonder that we sacrifice much to conventions and norms, social life would be ungraspably complex without them.

Bruce bn@bbn.com

Date: Mon Oct 26, 1992 2:00 pm PST
Subject: Re: Meanings and sentences

[from Gary Cziko 921026.1715 GMT] Eilieen Prince (921026) says:

> So,
>
> Give the book to me.
> Give me the book.
>
>may depend on what is being focused on and what is old information. The
>former, intonation being "normal" would mean that you should give the
>book to me, not to someone else. The latter would mean that what you
>should give me is the book, not someone else. (The interplay of
>intonation, context, and what has been said or done before makes this
>interpretation in fact more complex than I've just stated, but I think
>the idea is clear.)

Yes, I agree that they can be used quite differently. But then, why is it that I can say:

"Please contribute your money to my charity."

but not:

*"Please contribute my charity your money"

It seems that if I want to stress my charity over some other charity I would have to use emphasis such as:

"Please contribute your money to MY charity."

Why can't I use the dative contruction for "contribute"?--Gary

P.S. The fact that kids don't seem to make these and many other syntactic errors (although they make lots of semantic errors in figuring out the meanings of words) is taken by many linguists and language acquisition researchers as evidence that we are born with lots of innate, linguistic (especially syntactic) knowledge. I think this is wrong, but haven't figured out yet how PCT can help me support my belief. Maybe Avery can help.

Gary A. Cziko

Date: Mon Oct 26, 1992 2:15 pm PST
Subject: strategy

[From: Bruce Nevin (Mon 921026 08:24:02)] (Rick Marken (921023.0930)) --

>in the model the reference input "commands" a particular level of perception.

OK, I stand willingly corrected. But this somewhat "combative" distinction is not made, and the paragraph is less than clear in other respects too.

>People who are able to understand the paper (because they have
>already read PCT stuff) should see the implications just fine. Recall
>that I was writing the paper for people who had already passed PCT 101.

Then the reviewers are right if they say you are not really trying to get it published in this journal, i.e. that you are addressing a different audience. (That's my reading of the reviewer saying to Tom "you're determined to elicit rejection.") Nothing wrong with that if that's the aim, but there does seem to be a bit of conflict here.

  The fox, spying a bunch of grapes hanging above the path, leaps to pull it down but falls short again and again. Finally, with a last look up at the grapes the fox walks off, muttering "Probably sour anyway."

An imagined perception to counter the previous imagined perception, "mmm, how sweet and juicy those grapes will taste!"

If there really is no conflict, there will be no further efforts to publish in other than CSG publications, and no further complaints about reviewers and editors, right? :-)

(Gary Cziko 921024.0300 GMT) --

I agree with you that the "blind men" approach is a good one and ought to be pursued further. Had Einstein taken the approach to publication that some are advocating there might have been a small fringe group of relativistic physicists but they would probably have had little influence and would by now have all died off in the years following Albert's demise. (The fact that the parallel is obtuse, in that physics at least has a strong commitment to evidence, only strengthens the point.)

        Bruce   bn@bbn.com

Date: Mon Oct 26, 1992 4:45 pm PST
Subject: CALL FOR PAPERS & CALL FOR GUEST EDITORS -- S&G

```
  + + + + + + + + + + + + + + + + + + + + + + + + + + +
+ +                                                   + +
+                   Call for Papers                     +
+                          &                            +
+       Call for Guest Editorship of Theme Issues       +
+                         for                           +
+                 "Simulation & Gaming"                 +
+ +                                                   + +
  + + + + + + + + + + + + + + + + + + + + + + + + + + +
```

"Simulation & Gaming: An International Journal of Theory, Design, and Research" (S&G) publishes articles on simulation, computerized simulation, modeling, gaming, modeling, play, role-play, active and interactive learning, structured exercises and experiential learning. For 23 years, S&G has served as the leading international forum for the study and discussion of these methodologies as used in education, training and research.

S&G is a multi-disciplinary journal, covering a wide range of areas, such as sociology, political science, economics, psychology, education, environmental issues, international studies, management and business, policy and planning, decision making and conflict resolution, cognition, learning theory, communication, intergroup behavior, language training, media, educational technologies and computing.

Call for papers

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

The main types of contributions are (A) academic articles (empirical, conceptual, discussion, debate, descriptive, synthesis, etc. -- usually 3000 to 9000 words), (B) shorter articles (reports, communications, research notes, conference reports, accounts, instruments, etc. -- usually less than 2500 words) and (C) ready-to-use simulation/games (complete simulation/games with facilitator's instructions, participants' materials, debriefing schedule, etc. -- usually less than 3000 words).

Articles

Before submitting a manuscript, potential authors should write for a copy of the Guide for Authors, enclosing a self-addressed, sticky label and (in the USA only) $2 in stamps. Write to

David Crookall, Editor S&G
Morgan Hall, Box 870244
U of AL, Tuscaloosa, AL 35487, USA.

Simulation/games

Before submitting a ready-to-use simulation/game for ordinary issues of S&G, authors should write, enclosing a self-addressed, sticky label, and (if in the USA) $2 in stamps, to

Kuon Custer Hunt, S/G Section Editor
13735 NW Westside Road,
Yamhill, OR 97148, USA.

Call for Guest Editorship of Special Theme Issues
===

From time to time a special theme issue of S&G is prepared by a Guest Editor. Special issues in preparation or that have already appeared deal with business, debriefing, evaluation, military gaming and cross-cultural communication. In principle, any theme can be proposed for a special issue, as long it is important and of interest to a wide range of readers. If you would like to guest edit a special issue, please send:

- a one-page proposal (justifying the theme, outlining the rational, identifying possible authors and sub-topics, etc.)
- a short resume (one page)
- notes on any previous editorial experience
- name, address, telephone numbers and e-mail address(es) (the latter is essential)

to

crookall@ua1vm.bitnet or
crookall@ua1vm.ua.edu

To chat about it on the phone, call David at 205-348-9494 (w) or 205-752-0690 (h): please call back if not in -- do not leave a message.

Subscription inquiries about S&G should be directed to Sage Publications, PO Box 5084, Newbury Park, CA 91359, USA; telephone 805-499-0721; or 6 Bonhill Street, London EC2A 4PU, UK.

+ + + + + + + + + + + + + + + + + + + +

Date: Mon Oct 26, 1992 5:09 pm PST
Subject: various Re: presenting PCT

From Tom Bourbon [921026 15:40 CST] -------

[from Gary Cziko 921025.0200 GMT]

>To Bill Powers, Greg Williams & Tom Bourbon and other interested parties:

>When I try to discuss PCT with "mainstream" psychologists, two objections
>often come up: (a) feedback is too slow for many behaviors; (b) differentiated
>animals can behave with no sensory input. ...

>I remember Bill having made comments about the "speed" objection
>(and perhaps the deafferentation objection too) and Tom having made
>comments about the deafferentation studies and I would appreciate
>if they could summarize their arguments. From Greg I would expect
>some arguments that the lowest levels may not be fast enough for
>feedback control with the control effectuated by higher levels sending
>output commands and perceiving the results of the commands (although
>I don't want to start Bill vs. Greg on yet another topic before they've
>settled the current one).

>If these objections against PCT are so common, then perhaps PCTers
>should have objections ready against the objections.--Gary

I think Bill Powers [Bill Powers (921025.0800)] said most of it. His comments about the alleged slowness of feedback systems vs SR systems were right on target -- no surprise there!

As for deafferentation, I share Bill's opinion that the pursuit of the straw-issue of behavior without feedback has produced more mutilation and misery for hapless animals than nearly any other topic I can imagine - except maybe the vital research conducted by the cosmetics industry. Bill gave a good summary of many deficiencies in that research. He hinted at but did not mention directly the gross distortions of behavior that follow radical deafferentation. Even those researchers who seem the most determined to "prove" that behavior can occur after many sensory nerves are severed describe the loss of "elegance" in the actions of the victims. They understate the case! Following the all-important recovery period that Bill mentioned, the animals move "clumsily" and with obvious exaggerations that produce "feedback" via uninterrupted sensory paths -- there are ALWAYS such pathways. (If you like gore, read some of the studies in which researchers tried to eliminate as many sensory paths as possible -- first everything in the spinal cord, then many of the sensory roots of cranial nerves. Eventually, they decided that SOME sensory neurons must be left alone, or else the animal would die -- always.)

I am disappointed that people still find excuses to demonstrate the OLD "Law of Roots" -- the (general) functional difference between the dorsal and ventral roots of spinal nerves. There is no justification for re-doing that demonstration.

As for human clinical cases that resemble deafferentation, in his Principles of Psychology (1890), William James discussed the clinical work of a physician (in Germany I believe) with a young boy who had gross sensory deficiencies. James' account sounds remarkably similar to today's clinical accounts. Advocates of planned actions warp this literature severely to cite it as "proof" that feedback is not necessary for behavior. With extensive disruption of sensory pathways, humans act as "awkwardly" as experimental animals. The so-called awkwardness is often exaggerated action that produces sensations from remaining pathways. Close your eyes and slowly move your arm around. Try to become aware of as many sensations as you can, from as many places as you can. If you move vigorously enough, you will feel your arm movements all the way down to your toes. Imagine that you had a few months (years, in human clinical cases) to work with those remote sensations. Do you think you could "move, without sensory feedback?" Of course you could.

##

[From Rick Marken (921023.0930)]

Rick was replying to my post: Tom Bourbon (921022.13:45)

>For the benefit of voyeurs on this net, let me just explain what I think the

>> Carver-Scheier-Hyland-Lord-Hollenbeck.

>crowd are missing, and why people like Bandura are so hostile to PCT.

>According to PCT, behavior is controlled perceptual variables. In
>order to control those variables (relative to reference levels determined
>autonomously, by the organism itself) organisms must act to bring
>those perceptions to the reference level and counter any disturbances

>to the perception.

>The main goals of research in PCT are to determine 1) what perceptual
>variables an organism controls and 2) how this control is accomplished.
>The basic methodology for accomplishing 1) is "the test for the controlled
>variable" -- a methodology that is quite different than the traditional
>independent - dependent variable approach of conventional psychology
>(in fact, PCT shows that the conventional methodology reveals nothing
>about the nature of the organism; just statistical relationships
>between disturbances and compensating actions). The approach to accomplishing
>2) is modeling. Once we know what an organism might be controlling,
>we try to build working models that will control the variable in
>the same way. We consider the model a success when it acts just like
>the real organism (the "why 99%" criterion).

>So PCT goals and methods are quite different than those of conventional
>psychology. This makes conventional psychologists nervous and hostile
>because it is a major disturbance to their familiar way of going
>about their business. The Carver-Scheier-Hyland-Lord-etc crowd have
>made PCT acceptable by jettisoning the two aspects of PCT discussed
>above; they don't test for controlled variables and they don't model.
>They do IV-DV research and use statistics to determine if there was
>a significant effect of one variable on the other (with typically
>trashy results --such as the correlations reported by Tom). If these
>folks really understood PCT, they would know that they are wasting
>their time doing what they are doing -- but that would be true of
>other conventional psychologists too. People don't like to think
>that they are wasting their time -- so there is an easy solution -
>- don't understand PCT (in fact, don't even TRY -- hence, the hostility).

Right on the bulls eye, Rick. The people we are discussing nearly always start their articles, chapters and books with a brief (accurate) description of the simplest PCT model and an acknowledgement of Bill Powers. Then they identify the conventional areas in social and behavioral science to which they will compare control theory. Then, with bad IV-DV data in hand, they set about showing how a watered down control theory is compatible with everything conventional people embrace.

Do conventional people like positive feedback? No problem, we can put it in the control-theory model.

Do conventional people talk about behavior as a final product of cause and effect? No problem, we can say the control-theory model controls behavior instead of perception.

Do conventional people say behavior is controlled by its consequences or by plans? No problem .... and so on and so on.

No modeling. No awareness of why the glib changes are wrong -- yes, wrong, not merely questionable or aesthetically bothersome. It must be easy to publish that sort of thing; the literature of non-control theory is VERY LARGE.

###########################

Re the following from Chuck Tucker:
Date: Mon, 26 Oct 1992 08:48:44 EST
Subject: Statistics referred to in Tom's post of 921022.13:45

>   A TABLE SHOWING THE RELATIONSHIP BETWEEN
>   SEVERAL DESCRIPTIVE STATISTICS

Thanks, Chuck, for providing the entire table. For people who are new to CSG-L, or those who were here but did not save the table, it provides the quantitative support for my remarks about the badness of most statistically-significant results. Why 99%? Read the table!

Concerning your private post -- which paper? I access the mail from three different sites and at this one I cannot get to my archives.

####################################

[Gary Cziko 921024.0300 GMT]

>I have enjoyed the recent comments by Williams, Marken and Bourbon
>et al. on "selling PCT."

>But isn't there a third way to get PCT a wider audience which doesn't
>involve either (a) neutering and building shakey bridges to mainstream
>psychology or (b) presenting it as contrary to everything editors,
>reviewers, and readers have ever understood about psychology? I
>am thinking of (c) showing how PCT provides a unified theory of psychology
>in which the various flavors of psychology now existing (e.g., cognitive,
>S-R, reinforcement theory, etc.) can be seen as dealing with special,
>(very) limited subdomains of behavior and cognition. Yes, I am thinking
>of Marken's "The Blind Men and the Elephant" approach, dressed up
>to show how PCT can potentially handle the whole darn elephant.

>That is why I like Rick's paper so much and want to see this published
>very badly. But I am puzzled that I seem to be the only CSGnetter
>who is really excited about the potential of Rick's paper. Even
>Rick doesn't refer to it in his own comment on PCT (of course, we
>all know how modest and shy he really is)! Am I missing something
>wrong with Rick's unification approach? Or are other CSGnetters
>missing an important way of "selling" PCT?

Of course the third way is preferable. In our own way, Bill Powers and I do that in "Models and their worlds," the manuscript we have worked on since 1986 -- rejections began in 1989. We compare three models -- pure SR, pure plan-driven and simplest possible PCT -- in three different states of the environment. It is clear (to us, but not, so far, to reviewers and editors) that the PCT model is the general model that subsumes the other two. On reflection, many of our reviewers DID see that point -- they declared that is was so obvious a child would see it and that therefore we did not need to do experiments to show the point and our paper should not be published. I had forgotten their clear recognition of what we were trying to show.

Since 1988, the sequel to "Worlds" has been waiting in the wings, but we can't get the first paper published. In the sequel, the PCT model duplicates almost exactly the successes and failures of the other two models in "Worlds" -- all we do is change the reference signals in the model so that in one case it has a reference to sense its actions matching the sensed movements of the target in a tracking task, in the other it has a reference to perceive its actions matching a "memory" of what they were during an earlier run.

I believe participants in traditional IV-DV, C-E, S-R, I-O research do exactly those things: they control for perceptions of their own actions that match their understandings of what experimenters want to see. Consequently, researchers of every theoretical and methodological stripe are convinced that their "experimental hypotheses" are correct. They all miss the point. Of course, if "Worlds" is ever published, I am, certain we will face another six or seven years trying to publish the sequel.

Best wishes, Tom Bourbon

MEG Laboratory PAPANICOLAOU@UTMBEACH.BITNET
1528 Postoffice Street PAPANICOLAOU@BEACH.UTMB.EDU
Galveston, TX 77550 PHONE (409) 763-6325
USA FAX (409) 762-9961

Date: Mon Oct 26, 1992 5:34 pm PST
Subject: strategy

[From Rick Marken (921026.1500)] Bruce Nevin (Mon 921026 08:24:02)

Me
>>People who are able to understand the paper (because they have
>>already read PCT stuff) should see the implications just fine. Recall
>>that I was writing the paper for people who had already passed PCT 101.

>Then the reviewers are right if they say you are not really trying to
>get it published in this journal

Well, I don't think I'm not really trying -- but with the "Behavior of Perception" paper I admit that I am trying to get it published on my own terms (as a pro-PCT paper; not an anti-conventional psych paper). I've published plenty of the latter.

> The fox, spying a bunch of grapes hanging above the path, leaps to
> pull it down but falls short again and again. Finally, with a last
> look up at the grapes the fox walks off, muttering "Probably sour
> anyway."

I love it. I am like the "fox and the grapes" when it comes to "the blind men and the elephant". I think I'll find a away to fit Pygmalian and Galetea into my next paper; love those myths and fables.

Best regards   Rick

Date: Mon Oct 26, 1992 5:35 pm PST
Subject: Language perception and production

[Martin Taylor 921026 17:45]

A quick note on a couple of items of interest I encountered during my trip out West last week. Both have to do with language effects that the people studying them believe to be due to the perception of language, though they are manifest in performance. Neither researcher has come across PCT before, so far as I know.

(1) Sound shifts. John Ohala, U of Alberta, Edmonton.

As most people are aware, the sounds, words, and syntax of language change and drift over the generations. Ohala is studying the kind of sound shift that leads to pater-vater-father-pe're relations among words in a language family. Ohala has studied these relations in many different families all over the world. His thesis is that there are many ways of producing sounds that are perceptually very similar, and that a listener cannot always tell which of these mechanisms is used. A particular person will probably settle on one way of making a sound most of the time, but a listener trying to reproduce it may not use the same mechanism. Different minor articulatory errors make perceptually different sounds, so that, for example, a /p/ said with a slightly loose lip tension may not have a full closure, and a continuous fricative may be heard, like /f/. If a listener hears that as the correct form, he or she may articulate it as a lip-to-tooth closure rather than as a bilabial. Similar errors in voice onset time may lead to perception of /v/ instead of /f/. (These are my examples, not his, but look at what follows, on Elzbieta Slawinski's work).

As I understand him, Ohala claims that there is a pattern all around the world in which the sound shifts that are observed can be accounted for by shifts in articulatory mechanism between mechanisms that can have very similar perceptual effects when either is slightly affected by execution error. The acoustic percept is controlled and the desired sound is produced, but not necessarily by the same muscular actions as are used by the older generation. Particularly if there is explicit training ("Put your lips tight together and blow" as opposed to "Put your bottom lip near your front teeth and blow") and articulatory method could be propagated that differs across families or communities, like mutations in a genetic system. But it can be propagated without explicit training, by simple copying of the visual as well as the auditory perception.

The central point is that Ohala believes that it is the perception of the sound that is the constant across a sound shift, and that the articulatory mechanism leads to a shift in the range of error (which happens when control is not very tight--lowish gain), and this in turn leads to a shift in the central reference for the acoustic perception. When this latter has occurred in enough members of a community, the shifted percept becomes the norm, and someone who used the old articulation would be considered a foreigner.

(2) Changes of speech with age and hearing handicap. Elzbieta Slawinski, U of Calgary.

I did not get this one so clearly, but as I understand it, there is a shift with age in the transition rate of formants, and in voice onset time. Dr Slawinski believes that this is due to the perceptual ability of people to detect the transitions involved, and is trying to relate the speech of individuals to their ability to detect formant transitions and envelope transitions (voicing affects the amplitude envelope as well as the spectrum). She is considering very specific sound pairs, such as /ba/-/wa/ which differ in transition rates.

I gave Slawinski (and I think Ohala, but I'm not sure) the contact information for CSG-L. You may be hearing from them.

Martin

Date: Mon Oct 26, 1992 5:38 pm PST
Subject: Loose control

]From Jeff Hunter (921026-1)

> Bill Powers (921013.0930)
> > Greg Williams (921012-2) --

> >In principle, there is no difference between this sort of control
> >and the control of a cursor subject to a "hidden" disturbance -- in
> >both cases, what is tending to thwart control cannot be "seen."
>
> Qualitatively, perhaps not. But control is not just a qualitative
> matter -- control or no control. A control system that can cancel only
> 10 percent of a disturbance isn't much of a control system. In a
> tracking situation, 98 percent of the disturbance is cancelled.

Ok. Here's where we disagree strongly. You say, "A control system that can cancel only 10 percent of a disturbance isn't much of a control system." Here's an everyday counter-example.

During fall a squirrel can find lots of food. It stores (say) 30 grams of body fat as "insurance" for the winter. During winter it can find little food (except what it has hidden) and burns 0.3 grams per day. After 100 days winter had better be over. (Can you tell I'm Canadian?)

If the squirrel can reduce "fat burned today" by 10% (from .3 g to .27 g) it will stretch its fat reserves an extra 11 days. This could be the margin between life and death. (Now can you tell I'm Canadian? :-)

During winter the squirrel has great incentive to reduce "amount of fat burned today" to zero or even negative levels.

It presumably does this by fluffing its fur, sleeping during really cold patches, prospecting for nuts during warm spells, etc.

(We could test the reference levels of "fat burned" and "fat stored" by providing the squirrel with peanut butter, and finding how much the squirrel will eat when there is no cost or risk to gathering food.)

At any rate we have a control system (the squirrel) trying to control a disturbance (winter famine) in a CEV (fat burned per day) and not reducing the error (0.3 g) to zero, but still having a useful result (it survives till spring).

We get the same sort of thing in concert halls. The air conditioning cannot handle the heat output of 40,000 sweaty rock fans. However the management pre-cools the building, and keeps the conditioning running during and after the concert. Given that the hall is only used 4 hours out of every day the air conditioning (a classic control system) can keep the temperature acceptable even though it can only compensate for a fraction of the momentary disturbance.

> >Today my son Evan was having a problem with his new birthday
> >present, a radio- controlled truck. He asked me to help him figure
> >out what was wrong with the transmitter. Some experiments guided by
> >me showed a weak battery. Next time he'll be able to cure the

> >malady himself. No, he didn't hold a gun to MY head, either. We
> >BOTH got to where we wanted to be.
>
> See how easy it is when nobody is trying to figure out how to control
> someone else?

See why it seems that your definition of control is not:
- getting someone else to do what you want

but
- getting someone to do what they don't want to

> >This is the crux of our dispute. I claim that this is truly
> >CONTROLLING FOR, not just "wishes to see." B arranges A's
> >environment so as to encourage a class of actions by A which B
> >wants to see. If A doesn't perform actions in the class defined by
> >B, then B RE-arranges A's environment. And so on, until A does
> >actions in the class defined by B, or B gives up.
>
> If this is the crux of our dispute, then our dispute seems to come
> down to a quantitative question: loop gain. I guess I automatically
> dismiss examples in which the loop gain is so low that disturbances
> can't be significantly resisted. A model of the sort of situation you
> propose just above would, I imagine, have a loop gain very much less
> than -1; the degree of control possible would be very low. For
> significant control, I use a rough rule of thumb of a loop gain of at
> least -5 or -10. Only when the loop gain becomes that large do you
> begin to see the typical properties of a control system -- action
> opposing disturbance, controlled variable remaining near the reference
> level.

And I contend that you are blinding yourself to the possibility that loop-gain should be measured over a longer period if the disturbance is predictable (cyclic or following a known growth pattern).

In the case of Gary's son he is following a technique (involving yourself in your child's education) which has a high correlation with academic achievement. He may only be able to change Evan from a "B" to an "A" rather than an "A+" student. Nevertheless he is applying control, and it is useful.

... Jeff
De apibus semper dubitandum est - Winni Ille Pu

Date: Mon Oct 26, 1992 6:09 pm PST
Subject: Re: Language perception and production

Ohala would be a *very major* figure to get interested in PCT. E.g. if you were to ask people to name three top phoneticians I'm sure he would appear on almost everybody's lists.

Let's keep our cool. Avery.Andrews@anu.edu.au

Date: Mon Oct 26, 1992 6:25 pm PST
Subject: LIVING control system

from isaac n. kurtzer (921026.1743) to greg williams (not exclusively)

greg, since i'm assuming that you are the resident green thumb i'm asking you. PCT says that organisms (living) are control systems; does this apply to plants as well. like how flowers turn to the the light, or fold up at nighttime, or how a venous flytrap catches flies, etc. i realize that they have no comparable neural structures (actuall i am completely ignorant either way) but couldn't their interactions be modeled also?

i realize its a dial-a-cliche' question. i.n.kurtzer

Date: Mon Oct 26, 1992 6:49 pm PST
Subject: Bill Silvert, CSGINTRO.DOC, Verification

[From Dag Forssell (921026)]

For Bill Silvert, cc: Gary Cziko. General comment as inspired.

Bill, Greg & Pat have used my instructions below and tell me they got uue.c and uud.c without a problem. They have created an ASCII file which can be converted to uud.exe with the help of the DOS debug program. Instructions are included in the file. It was mailed to me today, and I will check it Wednesday or so, then forward it to you along with the revised starter.doc, which I propose to call csgintro.doc. I propose to call the new ascii file uudself.doc. Do you have a better suggestion? I propose a comment line for the csg/index below. Is that ok? I am anxious to have this csgintro.doc agree with what you put on the server. If it comes out different, it can certainly be revised.

I propose to put on the server this csgintro.doc, along with the uudself.doc, a revised and expanded (including illustrations) marken.doc (description of the spreadsheet model and, if you like, updated (current address, reformatted) demo1.doc and demo2.doc. I will mail disk to you and Gary both.

The csg/Index has plenty of room for notes on most files. Could you note with greater clarity which are binary and which ASCII?

Bill, I look forward to your comment.   Dag

enclosure: excerpt from revised CSGINTRO.DOC:

HOW TO OBTAIN TEXT AND PROGRAM FILES

A number of ASCII documents and binary computer programs are available on a fileserver maintained by Bill Silvert. It is possible to obtain all these files via e-mail. If you are on internet, it is easiest to obtain binary program files via anonymous FTP. If you are on MCI mail, you have read about how you can transfer binary files with Kermit or Zmodem protocols. (Type help at the MCI mail prompt for directions). But the server cannot send binary files to you over the internet mail network, so download uueself.doc so you can get the binary files uuencoded as ASCII files. The Internet address for the machine is BIOME.BIO.NS.CA. CSGnet files are kept in the subdirectory pub/csg.

To get basic information and a current listing of available documents, send a message as follows:

```
To: (Internet)         SERVER@BIOME.BIO.NS.CA.

Message:               help
                       ftp
                       get csg/Index
                       end

"help"            asks the server to send you commands and explanations.
"ftp"             requests the scoop on anonymous FTP for internet.
"get csg/Index"   requests the Index for the csg subdirectory.
```

In your message, pay attention to upper and lower case! DOS is not dos.

As part of the index (of the csg directory), you may be looking at:

```
programs/msdos:
dem1a.exe       128437  Bill Power's demonstration of perceptual control
dem2a.exe       123649  Bill Power's modeling of control

programs/source:
uudself.doc      54???  Compile uud.exe with DOS debug. Directions included.
```

If you want to request dem1a.exe (uuencoded) to get a "live" demonstration of the phenomenon of control, and the ASCII file uudself.doc which allows you to use DOS debug to create uud.exe to decode it, send the following message:

uue csg/programs/msdos/dem1a.exe
get csg/programs/source/uudself.doc
----
end of enclosure.

Date: Mon Oct 26, 1992 6:54 pm PST
Subject: Friendly PCT

[From Gary Cziko 921025.0050 GMT]   I said (921024.0300 GMT):

>>I am
>>thinking of (c) showing how PCT provides a unified theory of psychology in
>>which the various flavors of psychology now existing (e.g., cognitive, S-R,
>>reinforcement theory, etc.) can be seen as dealing with special, (very)
>>limited subdomains of behavior and cognition.

Rick Marken (921026.0930) replied:

>I do think this is a great idea. But there is still one little, teensie,
>weensie problem with this friendly approach; it is hard to keep your
>audience from noticing that their "flavor" of psychology is being revealed
>as a misinterpretation. The "Blind men" paper doesn't just show that the
>cognitive, S-R and reinforcement data can be seen as "special cases" of
>control; it also shows that the explanations of these data are wrong.
>For example, cognitive models (like the language models we've started
>discussing again on the net) are output generation models; the appearance
>(that cognitive behavior is generated output) is taken at face value and so
>the explanations have been completely off-base. So when you say that PCT
>is a way of integrating different approaches to understanding behavior,
>you are also saying that the conclusions that were based on that approach
>are wrong.

Maybe so, but don't you think it was first necessary to think of behavior as generated output and then see the problem with it? In this way, the generated-output view leads eventally to a better view and the generated-output people can be seen in a friendly way as leading to PCT. To use one of your favorite examples, do you think it was possible to come up with a heliocentric view of the solar system before coming up with a geocentric one? A generated-output view of behavior is probably better than one which says behavior is caused by angels or the soul interacting with the body through the pineal gland. But all these theories are just guesses anyway. It just happens that some of these guesses are easier to refute (S-R) than others (PCT) and so we think that PCT is closer to the truth (and probably is).

>there is no getting away from the fact that PCT says that
>psychologists can take their explanations of behavior and toss them in
>the waste basket. In terms of the "Blind men" analogy, it means that
>the fellow who has developed all those terrific models of the elephant
>that explain its "wallness" can now just throw those models away -- the
>elephant's "wallness" is just a side effect.

But if the blind man is never going to move from the side of the elephant, he might as stay with his "wallness" hypothesis. If I'm never going to venture beyond the Arctic or Antarctic circles, I might as well stick with my hypothesis that the sun rises and sets each day. And if a psychologist is going to stick with hungry rats and pigeons and schedules of reinforcement, he might as well stay with operant conditioning (as Greg Williams has said, Skinner was very good at working organisms, but not very good at understanding what makes them work).

>I think you have to take this into consideration, Gary, if you want to
>try to make the "Blind men" paper palatable. It's true that the
>feedback analysis integrates observations -- and I can see that that
>fact can be presented to psychologists in a friendly way. But it also
>shows (and there is just no getting around this) that current models
>of behavior (that were based on taking these observations at face
>value) are just, flat out, downright wrong.

But doesn't the "rightness" or "wrongness" depend on the context in one which is working? I don't think that NASA worries too much about the fact that the mass of the space shuttle approaches infinity as its velocity approaches the speed of light or that the engines increase their mass as they heat up. Newtonian physics works very well for NASA, thank you. But you can't stick with Newton if you want to understand physics in less limited domains, if you want a general theory.

>"So what if PCT let's me see SR, cognitive and reinforcment as special cases
>of a bigger picture; I'm only interested in reinforcement so I'll just keep
>doing my operant conditioning studies."

And you would have every right to do so. But people interested in a more general theory for understanding behavior will not have that option. We can give them another.

>I think the paper already gets
>this "so what" response; because it is only in a little section at the
>end that I gently suggest that this PCT point of view requires a whole
>new approach to explaining SR, cognitive and reinforcement phenomena --
>an approach based on recognition that what you are seeing is control of
>perception. So, I would really be interested in seeing what you have
>in mind to make PCT an attractive option for conventional psychologists.

But it's not really a whole new approach, is it? In some ways, its a welding together of cognitive and S-R ideas, as your paper shows. It's nothing but input-output connections through and through. Ah, but when you connect the output back to the input, include a reference level and increase the loop gain, something very new emerges--control. Emergence is trendy today, and we could push that as well. Control emerges from connecting input to output. Social phenomena emerge by putting individual control systems in a social setting.

Yes, I think the paper can be made friendly and influential as well. But it's just a guess. We won't know until we try.

Yours friendly, Gary.

Date: Tue Oct 27, 1992 5:51 am PST
Subject: Bandwidth

[From Bill Powers (921027.0500)]

Gary Cziko (921026) --

In case Martin Taylor's interesting treatise on bandwidth left some people scratching their heads:

Here's an example of the bandwidth of a control system. Hold up your forefinger about 18 inches in front of your nose and move it slowly from side to side over a total distance of three or four inches, like a slow metronome. Now, keeping the average position and the amplitude of movement the same, gradually speed up the movement, like a metronome going faster and faster. Keep going faster until you absolutely can't do it any faster. At that point you will be using your whole arm, and you will feel quite large muscular efforts, even though the movement from side to side is still only three or four inches (try to keep it that way).

The fastest movement you can produce is at a frequency essentially equal to the bandwidth of your finger position control system. Obviously you can perform this back-and-forth pattern at any slower speed (lower frequency) with no great difficulty, right down to zero frequency (stationary finger). But when you try to produce an oscillating movement at a frequency higher than the bandwidth, your control system simply won't obey. You can _imagine_ a faster movement, but you can't _produce_ a faster movement.

Why is there a bandwidth? One explanation might be that your muscles simply can't reverse the motion of your arm any faster, because they reach the limits of force that they can produce. If that were the only limit, you ought to be able to move your finger faster if you move it over a span of only a quarter of an inch instead of three to four inches. The maximum force needed to maintain an oscillation goes as the square of the frequency, so when you move your finger 1/10 as much, you should be able to oscillate your finger about three times as fast.

In fact, you can move perhaps a LITTLE faster, but certainly not three times as fast. You can oscillate your finger with an amplitude of, say, four inches or less at about 4 or 5 cycles per second, but not significantly faster, even for the smallest movements (I assume you're not a concert pianist, and anyway concert pianists don't have much occasion to practice sideways trills).

If you _increase_ the amplitude to a foot or eighteen inches, you will indeed find a decreasing speed limit set by muscle strength: the force required increases linearly with amplitude in a linear system (which your arm is not). At large amplitudes of movement you slow down because your muscles won't produce enough force to maintain the same frequency of oscillation you could maintain with a small amplitude. But below a certain amplitude, the speed limit is no longer set by muscle force. Something else is limiting the speed.

When you slowly speed up a small movement, keeping its amplitude the same, you'll notice another phenomenon. At low frequencies, you see a finger waving slowly back and forth. But at the highest frequency you can produce, you can see the finger only at the end of each movement where it reverses. Between those positions it's just a blur; you can see right through it. Obviously you couldn't track anything with your finger at that speed, because you couldn't see its movements, much less track irregular movements of something else. What you're seeing is the bandwidth of your visual perceptions of position. The frequency at which your finger just ceases to be a blur and becomes a finger again is the bandwidth of retinal position detection (actually you have to suppress eye movement by fixating on the background to find the true bandwidth, which is quite low, only 2 to 3 Hz).

It's interesting that the bandwidth or maximum frequency for small movements is higher than the bandwidth for retinal position detection. Something is limiting kinesthetic control at a frequency higher than that at which position control takes place, but at a lower frequency than is set by muscle strength. This probably involves a perceptual limit, too, in that kinesthetic position sensors do have speed limits, but more likely it is caused by temporal filtering that is required in order to make the kinesthetic control systems (that position your finger in the dark) STABLE.

The kinesthetic position control systems contain time delays of something like 50 milliseconds of neural transit time and synaptic delay around the loop. The muscles themselves have viscous damping. The noisy nature of neural signals, trains of impulses, requires that some smoothing take place in order to turn barrages of neural impulses into smooth changes in neurochemical concentration levels. All these factors mean that there is an unvoidable lag in these systems of about 100 milliseconds, part of it a transit-time delay and part of it an integrative or smoothing lag. That would imply that to switch as fast as possible from one position to another under kinesthetic control should take a little longer than 100 milliseconds, and to switch back another 100 milliseconds, for a total of 200 milliseconds for one cycle of a repetitive movement. That would give a frequency for continuous switching of 4 to 5 Hz, which is pretty close to what you see when you do it. Not bad for a ball-park estimate.

You can easily see the relationship between speed of movement and bandwidth. Try the experiment again, with small movements, only this time switch as fast as possible from one position to another 4 inches away, pause, then switch back as fast as possible, and pause. You're trying to generate a square wave. At low frequencies, each switch is discrete. You finger blurs over to the other position and is stationary for a while, then blurs back again. But as you increase the frequency of the square wave, still making each movement as fast as you can, the movements begin to blend into a continuous movement, so that when you reach the maximum frequency you're back to a continuous sine-wave movement. In fact, even at the low frequencies, each switch has been like half a cosine wave -- a high-frequency cosine wave at just about the bandwith frequency. So the slow square wave you started with was actually rounded off a little, and that rounding off meant that the movements actually never exceeded the maximum bandwidth for continuous oscillations.

It is possible for you to generate oscillations at higher frequencies. The only way to do it, however, is to destabilize your spinal control systems, the lowest level of control. If you press your hands together very hard and maintain the push until the muscles begin to fatigue, you may see "clonus" oscillations, at a frequency of about 8 to 10 Hz. This results from changing the force-tension curve in the muscles enough to make the control systems unstable. They break into spontaneous oscillation. But you can't produce this kind of frequency voluntarily. (You may see lower-frequency oscillations -- the next level may get unstable first. Shivering is probably a clonus oscillation of this kind, produced by destabilizing the control systems in some other way. So climb naked into the refrigerator if you want to see 10-Hz oscillations).

We're now approaching Rick Marken's territory. For visual tracking using control of finger position to follow a target, you obviously have to be able to see a finger while it's moving. This means that the bandwidth for following a randomly moving target is about 2 to 3 Hz, the frequency at which the finger just stops being a blur. This bandwidth is set by perception and output functions, not muscles. The kinesthetic systems clearly have a wider bandwidth; they can execute faster movements than you can control visually. And the lowest level of kinesthetic control, the spinal reflexes, have the widest bandwidth of all.

What's most interesting to me is that these nested bandwidths are just about what is necessary to maintain stable control at each level. There would be no point in being able to see movements beyond a bandwidth of 2 to 3 Hz because the kinesthetic control systems used by a visual-motor control system have a bandwidth only slightly higher -- 4 to 5 Hz. Therefore we DON'T see faster movements! In fact, if we could see faster movements, the bandwidth of the visual control systems would be so high that the lags of the lower control systems would be too long for stable control at the higher level. In technical terms, at a frequency where the phase shift of a sine-wave disturbance passing around the loop is 180 degrees, the gain would still be above 1. Negative feedback would turn into positive feedback at that frequency, and the whole system would oscillate. Oscillation is not good for control.

Rick Marken has explored several of the higher levels of perception, showing that as the (hypothetical) level increases, the bandwidth of perception continues to decrease. This is only logical, once you do some experiments yourself. For example, while moving your finger back and forth as fast as you can, _vary the amplitude_ between, say, a four-inch amplitude and a two-inch amplitude. Obviously, you can't even SEE "amplitude" in a time smaller than the fastest oscillation. And to VARY amplitude, you have to have a couple of oscillations of each size. In principle you could do one large oscillation and one small one, and so forth. In practice, you can't perceive changes in amplitude that fast. So you can't control amplitude as fast as you can control position. Rick's demonstrations are simple and elegant, as usual, showing the effect clearly. So naturally he can't get them published.

The relationships between bandwidths at different levels are, once you understand why they exist, perfectly simple and logical. It seems that bandwidth follows from physical principles and obvious relationships among physical phenomena, such as between frequency and amplitude. It's obvious that you can't change amplitude in less than one complete cycle, because amplitude doesn't even exist until at least one cycle is completed. Ho hum.

But remember that this is a constructed reality we're talking about. This relationship holds because of the way we perceive amplitude as a function of movements. Having constructed a perception of amplitude, we then discover that it has properties, and that it is related to lower levels of perception such as movement and position. The ho-hum self-evident relationship suddenly becomes evidence about how perception is constructed -- much more so than evidence about the natural universe. The bandwidth relationships also tell us that higher perceptions must be functions of lower ones, and that higher control systems use lower ones to accomplish their control. The evidence just continues to pile up that we are looking at -- and WITH -- a hierarchy of perceptual control systems.

When is the world going to wake up to what is going on here?

Date: Tue Oct 27, 1992 6:10 am PST
From: Bill Silvert
EMS: INTERNET / MCI ID: 376-5414
MBX: bill@biome.bio.ns.ca

TO: * Dag Forssell / MCI ID: 474-2580
CC: Gary A. Cziko
EMS: INTERNET / MCI ID: 376-5414

MBX: g-cziko@uiuc.edu
Subject: Your message with enclosures

>Bill, Greg & Pat have used my instructions below and tell me they got uue.c
>and uud.c without a problem. They have created an ASCII file which can be
>converted to uud.exe with the help of the DOS debug program. Instructions
>are included in the file. It was mailed to me today, and I will check it
>Wednesday or so, then forward it to you along with the revised starter.doc,
>which I propose to call csgintro.doc. I propose to call the new ascii file
>uudself.doc. Do you have a better suggestion? I propose a comment line for the
>csg/index below. Is that ok? I am anxious to have this csgintro.doc agree with
>what you put on the server. If it comes out different, it can certainly be
>revised.

I haven't checked the file, but if it works, fine. The instructions you encolsed seem fine also.

>I propose to put on the server this csgintro.doc, along with the uudself.doc,
>a revised and expanded (including illustrations) marken.doc (description of
>the spreadsheet model and, if you like, updated (current address, reformatted)
>demo1.doc and demo2.doc. I will mail disk to you and Gary both.

Why not send them electronically? Either ftp them or mail them.

>The csg/Index has plenty of room for notes on most files. Could you note with
>greater clarity which are binary and which ASCII?

No. The Index file is generated automatically every morning at 03:00 and it would take undue work to modify the creation program. There are about 50 types of files on this system which would have to be classified, not just ascii/binary, although currently this dichotomy is all that matters to the CSG users.
--
Bill Silvert at the Bedford Institute of Oceanography
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2
InterNet Address: bill@biome.bio.ns.ca

Date: Tue Oct 27, 1992 6:13 am PST
Subject: Up a level; language

[From Bill Powers (921027.0700)] Rick Marken (921026.0930) --

You're right that the Blind Men paper shows that the three Blind Men all have it wrong. They don't have it "partly right." How can you have part of a control system right? Unless you put all the parts together, you don't have a control system at all.

It's occurred to me that perhaps what we need is an "upalevel" paper. This should be written by you and Gary and perhaps other real psychologists. Instead of battering away at the wall, trying to get past the objections, why not start writing about the wall? "Look, folks, we have a new theoretical approach that has excited many respectable scientists from all over the world. But we can't get papers on it published in mainstream journals, simply because it contradicts accepted wisdom and because you have to understand it rather well before you can see what is new about it. We've offered papers on many subjects of interest to psychologists, trying to show how control theory can apply over the whole range of human behavior. But each time we do this, referees with narrow interests see it as wrong, misguided, outmoded, or old stuff. Referees with many different and even mutually-contradictory theoretical stances have said that their own field already takes control theory into account -- that not only is control theory compatible with their views, but they already know everything it has to offer. Or, at the other pole, the behaviorist says "You sound like a cognitivist," while the cognitivist says "You sound like a behaviorist."

"How can this be? Wouldn't you suspect that a theory that seems to mesh with widely differing schools of thought might have something of great generality and truth in it? But referees do not see it this way, because each of them judges from one narrow point of view. A theory that fits behaviorism as well as cognitive psychology might seem to offer a unifying principle of which these different schools are only special cases. But to see that, one has to see all the applications, not just one. One has to grasp the new theory as it is, not as it is imagined. How can we control theorists get past this barrier against publication so that psychologists in all fields can become aware of the potential of this approach?"

And so on.

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

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

I'm passing on social control for the time being.

Best to all,   Bill P.

Date: Tue Oct 27, 1992 8:32 am PST
Subject: Re: Upalevel

[From: Bruce Nevin (Tue 921027 10:23:52)]

Bravo, Bill! That sound just right to me.

I haven't yet been able to get my teeth into the current exchange on language. I have to think it through carefully, both intrinsically and also because I have not been attending to my evolving understanding of PCT re language and need to see just what state things have got to, simmering away there on the back burner. Soon, I hope, but BBN is keeping me busy.

Bruce   bn@bbn.com

Date: Tue Oct 27, 1992 2:31 pm PST
Subject: Three Worlds and Sequel

[from Gary Cziko 921027.0140 GMT]   Tom Bourbon [921026 15:40 CST]:

Thanks for your follow-up on the deafferentation question. The James reference is particularly interesting.

>Since 1988, the sequel to "Worlds" has been waiting in the wings,
>but we can't get the first paper published. In the sequel, the PCT
>model duplicates almost exactly the successes and failures of the
>other two models in "Worlds" -- all we do is change the reference
>signals in the model so that in one case it has a reference to sense
>its actions matching the sensed movements of the target in a tracking
>task, in the other it has a reference to perceive its actions matching
>a "memory" of what they were during an earlier run.

So why don't you guys combine the two papers into one? Wouldn't it be harder to use the "obvious" argument to reject it then?

In any case, I'd like to see a copy of the sequel if you are making it available.--Gary

Date: Tue Oct 27, 1992 3:47 pm PST
Subject: Little Man Version 2

[From Bill Powers (921027.1400)]

A paper by Bill Powers and Greg Williams, "A control-system model of human pointing behavior," has been submitted to Science magazine as a Research Report; it went off in the mail on Oct. 26. It is a review of the Little Man program, Version 2, and contains a comparison of motion trajectories with some human data from an article by Atkeson and Hollenbach.

The Little Man Version 2 is now available for distribution. It includes a 31-page writeup giving some details of the program, with block diagrams of the major aspects of the model.

The program includes two levels of kinesthetic control and one level of visual control. The kinesthetic control systems can be run in a test mode with square-wave reference signals at level 2 or level 1. All main parameters of the control systems are adjustable by the user. There are two modes of visual control: direct and mapped. In the mapped mode, a correction matrix is placed between the visual control outputs and the kinesthetic level 2 reference inputs. This map slowly adapts while the arm tracks the target (one operational mode makes the target jump randomly about in space). After an hour on a 486-33 computer, the map is essentially complete. After than, the trajectories of arm movement for large jumps of the target become very close to the way real arms move.

There is plenty of room for improvement at the visual control level. At the lower levels, Joe Lubin said "You've got it nailed." I hope that Joe will become active in carrying this model forward at Princeton.

The 360-K or 720-K disk contains full C source code with all necessary auxiliary modules and object files to compile under Turbo Pascal C version 2.0 (in addition to the runnable object code). The files are all contained in a self-extracting zipped file. There is an "install" batch file that creates a directory on drive C: with the necessary subdirectories and does all the necessary expansions. You can install from any floppy drive.

Prices:

For institutions and professional use:	$100
For members of the Control Systems Group:	$ 20

Send check and shipping address to:
William T. Powers
73 Ridge Place, CR 510
Durango, CO 81301

Specify 5-1/4 or 3-1/2 inch disk.

I encourage you to get your institution to buy the disk at full price. Note that if you join the CSG ($45 per year, $5 for students), you can get the disk for $20, at a total outlay considerably under the direct cost of the disk -- and you will also get 4 issues of Closed Loop. If you can't afford a copy for yourself, we will not object to copying the program from a legal owner. Strictly honor system, friends. Just remember that this is the sort of income that supports my work and Greg's. Don't ask me for a free copy unless you include a formatted disk in a self-addressed adequately stamped mailer and tell me a very sad story about how you live on even less income than Greg and I do.

Best to all,   Bill P.

Date: Tue Oct 27, 1992 4:32 pm PST
Subject: language

I have to to some other stuff before attending to the language thread I seem to have started, meanwhile thanks to Eileen for giving a substantive answer to Gary's questions about dative variants.

Avery.Andrews@anu.edu.au

Date: Tue Oct 27, 1992 5:28 pm PST
Subject: Re: language acquisition

(penni sibun 921026)

first, i think this discussion would benefit from not talking about sentences. sentences are not natural units of language. avery's initial discussion was not in terms of sentences, though he did use sentences as examples (quite reasonably: sentences are easier to talk about than more arbitrary-looking stretches of language).

second, i think avery's proposal is an interesting one. though, the subsequent discussion reads to me that the notion of Universal Grammar doesn't jibe very well with pct. maybe avery can comment on that.

third, i agree w/ what several people have already pointed out, viz., ``context'' is crucial in figuring out what a string of words means.

fourth, i'd like to make a brief argument for the interactionist point of view.

when people talk together, they are collaborating on what gets said. the person out of whose mouth some words or sounds come clearly has primary responsibility for what those words and sounds are, but, that person is not operating in a vacuum. just as an obvious instance of this, people often don't finish what they started to say, and often someone else finishes for them.

when i say something, i'm not taking a meaning in my head, encoding it, and flinging it at you, with you picking it up, decoding it, and putting it in your head. when i'm talking to you, i'm negotiating with you whose turn it is, i'm monitoring your reactions, i'm using words and phrases that i've learned are appropriate to this situation. the mechanism by which i learned that is the hard part, of course. how i decide what to say is the other hard part--or maybe it's not really. to sing the interactionist tune once more, i can lean on the world (including my interlocutors!) and my experience of it, so generally it's easy to know what to say next. my job as a linguist is to figure out what about the world makes it easy for me or anyone to use language. i don't think the answer will be found by focusing on sentences or other abstract grammatical structures.

fifth, as usual, bill makes the most insightful and interesting observations.

> [From Bill Powers (921025.1900)]
>
> To make a viable model out of this, it's necessary to specify the kinds of process that have to go on in the input and output functions. The input function must involve conversion from symbol structures into nonverbal experiences, so memory and association would be needed. The meaning-errors would have to be specified in terms of the dimensions of the meaning perceptions -- this is really a schematic of many control systems acting in parallel, each concerned with a different dimension of meaning. Then we have the problem of how to convert a difference between an intended meaning and a perceived meaning (in one of the dimensions) into an appropriate kind of change in the sentence -- that is, a change that will bring the perceived meaning of the new sentence closer to the reference-meaning.

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

> holes, doesn't it? And without the feedback loop, generation of meaningful sentences has to be treated separately from perception of the meanings in given sentences. The feedback model uses the same perceptual system in either case.

there's no empirical reason i know of either to treat them separately or not to.

> \>I'm thinking of the production system as a sort of transducer that
> \>turns 'meanings', some sort of internal structure, into an overt
> \>performance.

That's always seemed to me to be the weakness in the top-down kind of model.

i take avery to be saying something very simple, particularly by the word-choice ``transducer.'' on my view, there's very little difference between the internal form of what we are saying and our actually saying it; the difference is the same as that between whatever it is that leads up to kicking and a kicking action.

Given a meaning to express, there are very large numbers of sentences that will convey it. What kind of transducer can convert from a single signal into "the right" output among many outputs -- without any feedback to check if that was indeed the right output?

i don't understand this at all. if there are many right choices, why is it more difficult to determine if you've gotten *a* right one, than it would be to determine if you've gotten *the* right one?

I can see where exploring this model would require some novel experiments with language and meaning. It would be informative to see how people correct sentences that convey wrong meanings to them (even if the sentences are grammatically constructed). For a particular kind of meaning error, what sorts of changes in sentences do people make in order to correct them? This would begin to show the dimensions of meaning-error signals, and the dimensions along which sentences can be varied to alter particular kinds of meanings.
This takes a completely new approach to understanding how language and meaning hook together, I suppose.

the idea's not really novel at all (modulo the pct take on it), though it's not really to be found in linguistics. but many psycholinguists, conversation analysts, and at least one computational linguist have been asking these questions for a long time.

P.S. Given three nouns A, B, and C, you can create the sentences

The A B'd the C, and
The B A'd the C.

The captain hogged the ice-cream, and
The hog captained the ice-cream.

The only way to judge that one is allowable in ordinary discourse and the other is not is through the meanings of the words.

no, the only way to judge that either is allowable in ordinary discourse is to find it in ordinary discourse. if we played what-if games, we could create scenarios in which the first is not allowed and in which the second is.

[From Bill Powers (921027.0700)]

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

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

while i agree completely up to the semicolon, i don't think i can get my head around the bit that comes after. what if i see something neat and say, ``look!''. what's the relevant perception there?

The mere fact that we know of a sentence
shows that it's a perception. The same goes for grammar, for any regularity we PERCEIVE in language.

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

> The relationships between different levels of analysis of language are relationships among levels of perception, not levels of output production.

i don't think there are any useful levels of analysis of language, except possibly a distinction between the physical properties of language (eg, the sounds) and whatever else we do with it.

> We have to guess what the production processes are, because they're outputs and we don't perceive outputs.

this is interesting: we can never have any idea by what process we produce language??

> We perceive only their perceptual consequences: language is perception.

sure: language is perception; language is action; language is interaction. it's not sentences, though.

cheers. --penni

Date: Wed Oct 28, 1992 4:44 am PST
Subject: Botanical PCT

From Greg Williams (921028)

To whom it might concern: Address change for Willliam C. Littlewood, to P.O. Box 1171, Eastsound, WA 98245-1171.

Weather has been so nice lately that CLOSED LOOP will be delayed a few (I promise) more days so we can finish siding on another wall. Thanks for being patient.

>Isaac N. Kurtzer (921026.1743)

>PCT says that organisms(living) are control systems; does this apply to
>plants as well. like how flowers turn to the the light, or fold up at
>nighttime, or how a venous flytrap catches flies, etc.

At least some of those sorts of behaviors are controlling perceptions. Of course, this is most obvious when the effectors act reasonably quickly, and many plants don't have such quick-acting receptors (but still slooowly control some perceptions, like "roots down, shoots up," where "down" and "up" are perceptions based on sensing the direction of gravity. And plants' biochemical control systems are much like the biochemical control systems of other kinds of organisms. It is even beginning to look like plants have immune systems organized similarly to the immune systems of higher animals.

I realize that they have no comparable neural structures (actually I am completely ignorant either way) but couldn't their interactions be modeled also?

There is an article titled "The Secret Feelings of Plants" in the 17 October issue of NEW SCIENTIST (pp. 29-33) which reviews evidence that "nervous systems of plants and animals have more in common than was once thought." So far, no one has done PCT models of plant behavior; a ripe field for theses! If anyone gets serious about doing research in this area, I can provide additional references, going all the way back to Bose's studies on the sensitive plant -- I looked into all this as a neurophysiology student, for a seminar.

Best wishes, Greg

Date: Wed Oct 28, 1992 9:18 am PST
Subject: various Re: presenting PCT

From: Tom Bourbon [921028 9:07 CST] - To Gary Cziko [921927],

whose remarks were on the net, and to two others who commented off the net:

Concerning "Models and their worlds," you three suggested that Bill and I include the sequel, in which the PCT model emulates the performance of SR and plan-driven models. I believe the paper would be stronger with the emulations, and they provide a nice opportunity to build a genuine bridge to traditional behavioral scientists. It could follow this line: "Look at this, the PCT model can seem to act just like the other two. All it needs is a reference signal to perceive its actions matching either a perceived external 'stimulus' or a remembered 'plan.' Do you think people might be doing that in experiments where they KNOW they are EXPECTED to act like one of the traditional models, or where they KNOW they are only ALLOWED to act like one of the traditional systems? What do you think?"

But every reviewer who participated in a rejection of the paper criticized its length, as it is, without the emulations. You have probably seen this kind of anonymous remark:

"One would wish that the manuscript were considerably shorter."

"One would," indeed! Would one say such a thing were one's identity known?

What should we do? If we fail this time and try again we might add the emulations, but the manuscript will be even longer, and I am certain the bridge will be repulsed. If we had The Journal of Living Control Systems ... .

Best to all of you, Tom Bourbon

Date: Wed Oct 28, 1992 9:45 am PST
Subject: Squirrels, airconditioners, & linguistics

[From Bill Powers (921028.0630)] Jeff Hunter (921026-1) --

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

>During fall a squirrel can find lots of food. It stores (say) 30
>grams of body fat as "insurance" for the winter. During winter it
>can find little food (except what it has hidden) and burns 0.3
>grams per day. After 100 days winter had better be over.

>If the squirrel can reduce "fat burned today" by 10% (from
>.3 g to .27 g) it will stretch its fat reserves an extra 11 days.
>This could be the margin between life and death.

The importance of the result in the eyes of a biologist doesn't mean that there's a control system for the variable. In lower animals, many reference levels are inherited and are not adjustable. I would guess that the reference level for amount of stored fat (and stored nuts) is inherited -- the squirrel doesn't learn from experience how to adjust its own reference levels for body fat and stored food. So evolution quits when the inherited reference level is set high enough for most squirrels to survive most winters. Likewise for the mechanisms of hibernation; the squirrel's physiological systems turn down the reference signals for metabolism by an amount that's programmed in the genes. There's no way for the squirrel to say "Gee, if I just lowered my reference level for metabolism 10% more, I could survive a whole 11 days more." All that is out of the squirrel's control.

But the control system for achieving a certain amount of stored fat and stored food must be a pretty good one, especially if the margin for error is as small as you say. If the reference level for stored fat is 30 grams, and the loop gain of the associated control system is only 10, the system will stop with 10% error, or 27 grams of stored fat, and lose 11 days of reserves. If a disturbance that causes "fat burned today" to rise by 0.3 grams is poorly resisted, most of the 11 days will be knocked off the reserves.

>During winter the squirrel has great incentive to reduce
>"amount of fat burned today" to zero or even negative levels.

I doubt very much where an individual squirrel could (a) know anything about that incentive, or (b) do anything about it if it knew. When a mountain climber falls into a crevasse, there is a great incentive to learn to fly. That doesn't make it possible.

>At any rate we have a control system (the squirrel) trying to
>control a disturbance (winter famine) in a CEV (fat burned per day)
>and not reducing the error (0.3 g) to zero, but still having a
>useful result (it survives till spring).

You'd better draw me a diagram of this control system. What variable does it sense, and how? How does the error get turned into an action that opposes the effect of a disturbance on the variable? I suspect that we have something here like "biologist's purpose." You know, squirrels store fat in the fall in order to survive the winter. The implication is that the squirrel senses survival, and if survival falls below the reference level, increases its fat-storing and nut- gathering to bring survival back up to the reference level. Of course that's pretty hard to do for a squirrel that detects less than 100% survival.

Not all outcomes are controlled.

>We get the same sort of thing in concert halls. The air
>conditioning cannot handle the heat output of 40,000 sweaty rock
>fans. However the management pre-cools the building, and keeps the
>conditioning running during and after the concert. Given that the
>hall is only used 4 hours out of every day the air conditioning (a
>classic control system) can keep the temperature acceptable even
>though it can only compensate for a fraction of the momentary
>disturbance.

The only reason this works is that the fans are willing to accept a wide range of temperature error, beginning at "too cold" and ending up at "too hot". In fact, during the concert there is no control of temperature at all: the heat input is larger than the ability of the cooling system to resist. The air conditioner runs continuously at maximum cooling -- the output doesn't vary with temperature error, so the loop gain is zero. This is not only a lousy control system, it isn't a control system at all. Management is essentially just turning on the refrigeration and hoping there aren't too many complaints. That's not very good control by management, either, assuming that management even cares about the complaints.

The only control that exists is over a 24-hour period. The average temperature is kept within plus or minus x degrees of a reference level. That's not very good control. But maybe it's good enough for the controller. When the loop gain is low, we can assume that it isn't very important to control that variable at exactly the reference level. Not important enough to buy a bigger air conditioner.

A good control system keeps its error to a small percentage of the reference level. You can't make a poor control system into a good one by pointing out that it's good enough for a not-every-important purpose (important in the eyes of the controller, that is -- not the observer, who may know more than the controller does). You can't even make it a good on by showing that it serves an important purpose. If it allows large errors, it's not a good control system.

RE: Greg helping Evan to solve a problem, at Evan's request:

>> See how easy it is when nobody is trying to figure out how to
>>control someone else?

>See why it seems that your definition of control is not:
>   - getting someone else to do what you want
>but
>   - getting someone to do what they don't want to

It isn't so much getting someone to do what they don't want to as getting someone to do what you want, without considering what they want. You're controlling in either case, of course. The main difference is whether your attempt at controlling what matters to you results in resistance from the other person. If the other person has asked you to help, you can be pretty sure that complying will not result in resistance and that you're helping with something that the other person wants to be helped with.

(Penni Sibun 921026) --

>first, i think this discussion would benefit from not talking about
>sentences. sentences are not natural units of language.

Gee, they seem like pretty natural units while I'm writing. You seem to write in well-formed sentences most of the time. Is that just an accident? Or is it something you control for?

>third, i agree w/ what several people have already pointed out,
>viz.,``context'' is crucial in figuring out what a string of words means.

I get uncomfortable when people talk about what words mean, at least when the discussion is theoretical. Words don't mean anything; they're just noises or marks. Meanings arise when people hear or read words. When we hear someone else speaking we experience meanings, which I suppose to be the (mostly) non-verbal experiences that the words evoke through association with memory. We may also be puzzled by such meanings, and wonder what the words meant to the person who spoke them -- sometimes it's obvious that they couldn't mean to the other what they mean to me. Then we try to work this out in conversation, using different words, gestures, intonations, and so forth to evoke still more meanings, until we finally find a set of meanings to which neither of us objects. I suppose that's what you mean, in part, by context.

>when people talk together, they are collaborating on what gets
>said. the person out of whose mouth some words or sounds come
>clearly has primary responsibility for what those words and sounds
>are, but, that person is not operating in a vacuum. just as an
>obvious instance of this, people often don't finish what they
>started to say, and often someone else finishes for them.

Sounds OK to me. In group speech you can interrupt, and you're in face-to-face contact. The meanings are building up in each person's head as the conversation goes back and forth, providing a sense of what the conversation is about (other than words). Sometimes it's just a sense of pleasantness, like "This is nice, we're having a good talk." Sometimes there's more substance to it, as when we're haggling over a price or trying to arrive at a mutually-agreeable way of expressing an understanding of something (and trying to find out just what the heck the other person thinks he or she understands by my words).

>when i say something, i'm not taking a meaning in my head, encoding
>it, and flinging it at you, with you picking it up, decoding it,
>and putting it in your head.

I agree that meanings aren't "encoded" into words and then "decoded" again. Meanings exist in the person hearing the words; the words are pointers to them. Nothing but words, encoded as sound-waves, passes between people. They carry no extra baggage of meaning with them. We can't "transmit" meanings. We can only evoke them.

But the other meaning of your statement depends, I think, on the kind of conversation we're having. If I'm trying to get across to you some well-formed meaning in my head, I choose words that I hope will evoke some similar meaning in your head, which you get by drawing on your own nonverbal experiences. You may then try describing the meaning you got by using different words and examples, or you may tell me or show me what you would do on the basis of the meaning you got from my words. Then I take the meanings I get from your words, gestures, and so forth and compare them with the meaning I'm hoping to convey, and if there are differences, try to express them in ways that will convey the difference; or (if I don't understand why there's such a difference) I may try a different example or way of describing my meaning, hoping to perceive less error the next time around.

Of course in casual non-purposive conversation, the conveyance of meanings and checking up on what was conveyed is not the main point. The participants just assume that everyone gets the meanings that were intended, because the exact meanings aren't very important and the conversation doesn't have any particular agreed-on goal. The main goal is just to be in a conversation.

>to sing the interactionist tune once more, i can lean on the

>world (including my interlocutors!) and my experience of it, so
>generally it's easy to know what to say next.

That's a little too casual for my taste. I don't buy this "leaning on the world" bit, because it implies that the world is the same for everyone, and that its properties and behaviorial characteristics don't require any perceptual interpretation to be known. When I'm not being a theoretician, of course, I take the world as it appears to be and don't constantly remind myself of all the interpreting my brain must be doing to make the world look that way. But here I am being a theoretician, so I have to worry about such things. The naive-realist approach to the world doesn't cut it, theoretically.

>my job as a linguist is to figure out what about the world makes it easy
>for me or anyone to use language. i don't think the answer will be found
>by focusing on sentences or other abstract grammatical structures.

What if it isn't something about the world, but something about the brain's way of perceiving a world that makes language seem easy to use? If it were the world that made language easy, dogs would speak English (in this country). I agree with you that the answer isn't to be found in analysis of sentences or abstract grammars. But evidence is to be found there, hints about the way the brain deals with its experiences and symbolizes them.

Does being a linguist mean ignoring questions about how the brain works and what the brain has to do with language and experience? Does interactionism confine you strictly to what you can observe passing between people, to the surface appearances of their behavior and their interactions?

>>The input function must involve conversion from symbol structures
>>into nonverbal experiences ...

>where are the symbol structures coming from?

They are created by the way we perceive streams of incoming or self- generated words. We inmpose certain structurings on the stream -- sequence, betweenness, associations between words. These structurings, as well as the individual words, come to evoke meanings for us. When we generate communications, we do it by creating similar structurings that evoke in ourselves the meanings we want to communicate; by using our articulation machinery, we create streams of words which we perceive just as we would if they were generated by someone else. As we do so, the same streams are caused to enter someone else's perceptual system, where similar processes evoke meanings at many levels of perception.

>what does meaning have to do with anything? why can't speech actions be
>just like other actions, which presumably aren't mediated by ``meaning''?

shprtnof frmtz quarntv sprtstx eermoffs wr smorfpit chlathrogh.

That's what meaning has to do with anything. It has EVERYTHING to do with language.

Speech ACTIONS are like any other ACTIONS. They consist of varying muscle tensions that move the physical equipment around. Speech PERCEPTIONS, on the other hand, are the controlled results of those actions. Speech perceptions, at the lowest levels, provide the perception of noises, which are perceived as words, which are perceived as word structures, which are perceived as instances of grammatical forms, and so on (depending what you believe about higher structures). At the same time, the speech perceptions at each level evoke nonverbal meanings -- remembered experiences at various levels, which become an imagined picture that is built up as the words and structures continue to appear at various levels in the perceptual system. That imagined picture is the meaning of the communication being perceived. The meaning is simply a nonverbal experience, like the one you get from looking silently around at the world, or feeling it or tasting it and so on.

>on my view, there's very little difference between the internal form of
>what we are saying and our actually saying it; the difference is the same as
>that between whatever it is that leads up to kicking and a kicking action.

Our "actually saying" something amounts to a perception of ourselves saying something. There's no other way to know you're saying something. Your experience of your own saying is a PERCEPTION, not an ACTION. The action of saying something consists of sending signals into muscles, causing them to contract. You don't experience that part. All you experience are the tactile, kinesthetic, and auditory consequences of having created those outputs. Others, of course, also experience consequences of those outputs. Their experiences of them are similar to, but not identical to, yours. They may _hear_ "hello" much as you do, but they won't know how it feels to you to _say_ "hello."

The same goes for kicking. You feel your leg movement, you feel the contact with the kickee (if any), you may see your leg move. But you don't experience the signals going down your spine to the muscles, and you don't experience the contraction of the muscle fibers. All you can know about kicking, someone else's kicking or your own, is what you perceive. Perception is input, not output. All you know about your own behavior is in the form of perception, which is information coming into you, not going out of you.

\>>What kind of transducer can convert from a single signal into "the
\>>right" output among many outputs -- without any feedback to check
\>>if that was indeed the right output?

\>i don't understand this at all. if there are many right choices,
\>why is it more difficult to determine if you've gotten *a* right
\>one, than it would be to determine if you've gotten *the* right one?

You can't "determine" anything unless you perceive it. Perception is not associated with outgoing channels in the brain; only with incoming, afferent channels. So a transducer in the output channel will have effects, but they can't be known to the system. They are blind actions. The system employing the transducer can't know the results of the transduction unless those results are represented in perception, which is an input to the higher system fed back from the consequences of the action. All that the higher system can know about the results of transduction is contained in those perceptions; all that can be controlled is in those perceptions. That's what PCT is about: control of perceptions.

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

\>while i agree completely up to the semicolon, i don't think i can
\>get my head around the bit that comes after. what if i see
\>something neat and say, ``look!''. what's the relevant perception
\>there?

What you see is a perception. Your feeling that it is neat is a perception of the perception of it. Your experience of saying "look!" is a perception. Your experience of the other person who then looks is a perception. It's ALL perception. A sentence is a perception simply because you are aware of it: everything you are aware of is a perception. If you construct a sentence, what you have constructed is a perception, not a thing outside you. If you are aware of having a thought about the sentence, the thought is a (higher-level) perception. The whole world we experience, whether concrete or abstract, is made of perceptions in the afferent parts of the brain at various levels of organization. Understanding this is the sine qua non of understanding PERCEPTUAL control theory. Control systems control their inputs, not their outputs.

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

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

\>>   The relationships between
\>>   different levels of analysis of language are relationships

>>   among levels of perception, not levels of output production.

>i don't think there are any useful levels of analysis of language,
>except possibly a distinction between the physical properties of
>language (eg, the sounds) and whatever else we do with it.

You're bluffing. You don't know what levels of analysis may exist, or whether they might be useful. You're just saying that you don't intend to pay any attention to levels of analysis, whether they're useful or not.

>> We have to guess what the production processes are, because
>> they're outputs and we don't perceive outputs.

>this is interesting: we can never have any idea by what process we
>produce language??

We can guess, but we don't experience our own production processes. All we experience are the perceptual consequences of them, such as feeling our tongues and jaws move and hearing the sounds we make.

>> We perceive only their perceptual consequences: language is perception.

>sure: language is perception; language is action; language is interaction.
>it's not sentences, though.

Not quite. Action, as we know it, is perception, and so is interaction. All we know is perception. Language is not sentences in the same sense that cars are not wheels. But in any case, we perceive sentences, wheels, and cars. How else could we know about them?

Am I getting across at all?

Best to all,   Bill P.

Date: Thu Oct 29, 1992 11:59 am PST
Subject: cat experts

David Goldstein 10/29/92

My son Joshua made an interesting observation about our one-year old cat named Chelsea. Josh noticed that when he turned a ratchet wrench, Chelsea started to lick. When he stopped, Chelsea would stop licking. This was a reliable phenomenon, at least on the day when Josh discovered it. We told him to stop "making Chelsea lick." after a while.

Can any cat experts out there explain this?

As far as I know, Chelsea was never classically conditioned to lick to the sound of a ratchet wrench.

I think it is the sound of the ratchet wrench. He stopped licking when Josh stopped making the sound.

What is the perceptual variable which is being disturbed by the sound of the ratchet wrench which the licking is correcting?

Does any body else's cat do this?

Date: Thu Oct 29, 1992 3:36 pm PST
Subject: Re: Squirrels, airconditioners, & linguistics

(penni sibun 921029)

my injuries are acting up, so my replies will be succint.

   [From Bill Powers (921028.0630)]

   (Penni Sibun 921026) --

>first, i think this discussion would benefit from not talking about
>sentences. sentences are not natural units of language.

Gee, they seem like pretty natural units while I'm writing. You seem
to write in well-formed sentences most of the time. Is that just an
accident? Or is it something you control for?

though we intellectual academics tend to forget, most language use is not writing. since illiterate people are competent language users, a theory of language use clearly should not depend on a theory of writing.

>third, i agree w/ what several people have already pointed out,
>viz.,``context'' is crucial in figuring out what a string of words
>means.

I get uncomfortable when people talk about what words mean, at least
when the discussion is theoretical. Words don't mean anything; they're
just noises or marks. Meanings arise when people hear or read
words.

i agree words don't have meanings. put ``string of words'' in scare quotes.

>when i say something, i'm not taking a meaning in my head, encoding
>it, and flinging it at you, with you picking it up, decoding it,
>and putting it in your head.

I agree that meanings aren't "encoded" into words and then "decoded"
again. Meanings exist in the person hearing the words; the words are
pointers to them.

what is the difference bet. an encoding and a pointer?

Nothing but words, encoded as sound-waves, passes
between people.

plus a whole lota other ``contexty'' stuff.

But the other meaning of your statement depends, I think, on the kind
of conversation we're having. If I'm trying to get across to you some
well-formed meaning in my head,

most negotiation in language is not explicated or verbal. negotiation of turns, for instance, is going on all the time.

>to sing the interactionist tune once more, i can lean on the
>world (including my interlocutors!) and my experience of it, so
>generally it's easy to know what to say next.

That's a little too casual for my taste. I don't buy this "leaning on
the world" bit, because it implies that the world is the same for
everyone, and that its properties and behaviorial characteristics
don't require any perceptual interpretation to be known.

why does it imply that? i don't think it does. leaning on the world is an interaction bet. an individual and (the individual's perception of) the world.

Does being a linguist mean ignoring questions about how the brain
works and what the brain has to do with language and experience?

no. there's lotsa cool neurolinguistic work going on (and if i had a different background i might be doing some of it). the brain's part of the world too.

Does
interactionism confine you strictly to what you can observe passing
between people, to the surface appearances of their behavior and their
interactions?

no. what it confines you to is analyses that are not restricted to organisms in
isolation.

>where are the symbol structures coming from?

They are created by the way we perceive streams of incoming or self-
generated words.

this assumes perception of words, which may be a faulty assumption.

We inmpose certain structurings on the stream --
sequence, betweenness, associations between words.

perhaps we don't in fact decompose it.

These structurings,
as well as the individual words, come to evoke meanings for us. When
we generate communications, we do it by creating similar structurings
that evoke in ourselves the meanings we want to communicate; by using

why assume we *create* language structures rather than reuse ones we know (eg, by having
heard them)?

Speech ACTIONS are like any other ACTIONS. They consist of varying
muscle tensions that move the physical equipment around. Speech
PERCEPTIONS, on the other hand, are the controlled results of those
actions. Speech perceptions, at the lowest levels, provide the
perception of noises, which are perceived as words,

it's not clear that noises are necessarily perceived as words or that the noises are
sufficient for perceiving words. why are you so certain that all this analysis is going
on? to me, your analysis here sounds like the sort of thing that you criticise by saying
``so and so is taking the perspective of someone sitting up in the sky and looking down
at the world.''

>on my view, there's very little difference between the internal >form
of what we are saying and our actually saying it; the >difference is
the same as that between whatever it is that leads up >to kicking and
a kicking action.

Our "actually saying" something amounts to a perception of ourselves
saying something. There's no other way to know you're saying

The same goes for kicking. You feel your leg movement, you feel the

we seem to be agreeing here.

>>What kind of transducer can convert from a single signal into "the
>>right" output among many outputs -- without any feedback to check
>>if that was indeed the right output?

>i don't understand this at all. if there are many right choices,
>why is it more difficult to determine if you've gotten *a* right
>one, than it would be to determine if you've gotten *the* right
>one?

You can't "determine" anything unless you perceive it. Perception is
not associated with outgoing channels in the brain; only with
incoming, afferent channels. So a transducer in the output channel
will have effects, but they can't be known to the system. They are
blind actions. The system employing the transducer can't know the
results of the transduction unless those results are represented in
perception, which is an input to the higher system fed back from the
consequences of the action. All that the higher system can know about
the results of transduction is contained in those perceptions; all
that can be controlled is in those perceptions. That's what PCT is
about: control of perceptions.

we must be at crosspurposes here. i can't find your original text, but at some point i believe you had said that it was harder for a transducer to produce one of several ``correct'' outputs than a single ``correct'' output, and i've been trying to figure out what you meant.

I agree in the sense that we use words in any way needed to evoke meanings in ourselves that we hope are similarly evoked in others. But I disagree with your extreme position, because it's clear that people DO make well-formed sentences and perceive them as such; you did so in

most language users have no idea what a well-formed sentence is; ``sentence'' is not a category for them. it's irrelevant that you and i can make observations about sentences in a metatheory that includes the concept of sentence.

the sentence above. We have concerns not only with conveying meanings, but with observing grammatical rules and customs.

and we have concerns about not spitting and not shouting and being polite and and....

Syntax and grammar are aids in narrowing down possible meanings

so are drawing syntax trees and encyclopedia searches.

we control for them the most carefully when we are most concerned with getting an exact meaning across.

if we are hyper-educated and have devoted considerable energy to learning the skill.

once you start *reflecting* on grammar, etc., it's at a ``higher level of perception'' as you put it, and not part of the process of producing language, which, you have argued, is transparent.

>i don't think there are any useful levels of analysis of language,
>except possibly a distinction between the physical properties of
>language (eg, the sounds) and whatever else we do with it.

You're bluffing. You don't know what levels of analysis may exist, or whether they might be useful. You're just saying that you don't intend to pay any attention to levels of analysis, whether they're useful or not.

i'm being extreme. i believe none of the reflective levels we've been discussing is going to help us figure out what is really going on. we need to dump them and look for new levels (if any).

>sure: language is perception; language is action; language is
>interaction. it's not sentences, though.

Not quite. Action, as we know it, is perception, and so is interaction. All we know is perception.

sure, in the end, everything is perception. but a theory needs to be a little richer than that to explain things. and i feel interaction is a major enriching requirement.

Language is not sentences in the same sense that cars are not wheels.

no, language is not sentences in the same sense that space is not ether.

cheers. --penni

Date: Thu Oct 29, 1992 9:21 pm PST
Subject: Language

RM72/ [From Bill Powers (921029.2200)] Penni Sibun (921029) --

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

I agree, language theory has to handle spoken language. It may be that written language is somewhat artificial -- but if you look at transcripts of spoken language, you have to wonder how much imagination the listeners have to use to get a meaning out of it, much less the intended meaning. On the other hand, I do agree with your general idea, which is that language involves more than one person, and involves an interplay between speakers/writers, not just mechanical translations of rules into outputs and back into inputs.

I think I want to go even farther than you do when you say that words don't have meanings but strings of words do. Today I finished reading Nicolas Freeling's latest mystery, Flanders Sky. Freeling writes in a style like none other; very dense, hard to read because you can't skim or even skip if you want to get the point. A slow but satisfying read. What came through to me today was a fresh appreciation of the fact that words are not what they point to.

Freeling writes about people who know several languages, and makes references, jokes, allusions to sayings with lots of German, French, Flemish, and English in them. At one point the word "but" came up somehow, and it struck me that "but" does not itself have a meaning but points to a perception of "butness" that is experienced by people who speak all different languages. We have "but", "aber", "mais," and (from Latin) "sed," all of which are just arbitrary gabble, but all of which are understood by speakers of different languages to mean a perception of a particular kind. The meaning is a sort of a pause, a contradiction in the offing, a turn in the line of argument from plus to minus or at least in a different direction -- it takes a lot of words to refer to the meaning without using "but" or an equivalent, because the meaning is not a word, but a perception. It's a high- level perception of some sort; I don't know what sort. People agree on the perception enough to allow translating easily from one word for it into another. Maybe there are different nuances, but the basic experience is the same.

What came through more clearly today was the sense of language having two parts, the least important one being the most obvious, the words, phrases, sentences, and so forth. The most important part is the hardest to put into words and the least obvious. It is the background of images, associations, relationships, reactions, intentions, and imaginings that flows and interacts just beneath the surface while the mechanical aspects of language, the tokens and structures or whatever you want to call them, march across the field of consciousness.

This underlying field of nonverbal perceptions is made of memories of experiences. I tried to express this concept to Bruce Nevin half a year ago or more, but it's very hard to express, much less express clearly. This idea gets mixed up with the fact that language itself is an experience -- if you don't have any meanings hooked up to the word-flow (as when you don't know the language you're hearing), it's clearly a nonverbal experience. It's not so clear that it's still a nonverbal experience even when you do know what the words mean. You can learn categories, sequences, and rules in such a language without ever knowing what the words mean, just as you can learn "i before e except after c" without knowing what the word means: carmotteis is probably misspelled if it's an English word.

This is how a computer learns to use words. You don't actually have to use real words in any computer program for parsing language. You could number the words in dictionaries and construct sentences using just the numbers. The rules of language would apply directly to the numbers. If 1289 is always followed by a number from the set {2365,12, 99256,...} you have a "linguistic rule." If you learn this rule by heart, you can construct valid expressions even not having the dictionary around and thus having no hint as to what you are actually saying. A Frenchman might think that "Keepen der fingerpoken offen der blinkenlights" is really German. If someone had encoded protocols in numbers and had given them to Harris, he could still have done his analysis, or at least the part of it involving the determinations of transition frequencies.

What's missing from a sentence or other communication made of the key numbers instead of the words is meaning -- the most important part. I was trying to suggest to Bruce that while there are certainly purely linguistic rules, rules that can be applied without necessarily knowing what the words mean, there are much more important rules. Those rules are the rules of the world of experience. You can say "The boy balanced the orange sunset" without violating any rules of grammar or syntax, but the usage factor for such a sentence would have to be extremely low. That's not because of language, but because it's hard even to imagine a scene that would go with the sentence. It's not language, but experience that tells us that sunsets aren't the sort of thing you balance. Of course there's a great urge to make up a scene that would explain such a sentence -- the boy is a painter! Now we expect the next sentence or an addition to the existing one to tell us what he balanced it against -- a reflection from a window, and so on. Only when we can imagine something that makes experiential sense do we say that we made sense of the sentence that evoked the bits and pieces of the experience. Language is only the means by which we evoke memories of experiences in other people. That's the thesis I'm pushing. And the converse must be that behind language there must be a great organizing force that comes from our nonverbal experience of how the world works.

The difference between an encoding and a pointer, since you asked, is that the concepts of encoding and decoding imply some objective rule and its inverse. Meaning is not transmitted that way. My meaning is put into shape for communication by finding ways of talking that evoke the right meaning in me. The listener, hearing a communication but not its meanings, experiences non-verbal perceptions brought out of the listener's experiences, not the speaker's. The encoding and decoding processes, if you want to call them that, are neither objective nor related as inverses. They depend on the experiences of the speaker and those of the listener, and on private associations each has formed between words and nonverbal experiences. Formalisms of language may facilitate this process, but can never make it perfect or (I suspect) even near to perfect. As our very network demonstrates, the sender's meanings do not automatically pop up in the receiver. What the receiver gets from the words is often far from what the sender hoped to be posting. Meanings are not "received." They are constructed out of the experiences of the hearer.

>we must be at crosspurposes here. i can't find your original text,
>but at some point i believe you had said that it was harder for a
>transducer to produce one of several ``correct'' outputs than a
>single ``correct'' output, and i've been trying to figure out what you meant.

Not quite what I meant, as I construct a meaning from your words. A blind transducer can be commanded to produce outputs. But there is no way for the commanding entity to know what those outputs actually are unless they are monitored -- perceived. The commanding entity can't know whether any output is the right one OR the wrong one without direct feedback from the result. And then it's really the fed-back result that is being adjusted; that can be several steps removed from the actual output.

The standard way around this, on which all top-down models rely, is to assume that SOMEHOW the transducer evolves so that it only puts out right outputs. Since the right output may depend very strongly on the state of the outside world, this SOMEHOW becomes a tremendous stumbling-block to making a working model.

This is the unfortunate legacy of the computer revolution. Computers, unlike nervous systems, are designed to do exactly what they are told every single time, with no errors. Even their outputs are designed to be error-free: if the computer commands that an "A" be printed, an "A" always appears unless the computer is physically malfunctioning, in which case you can't trust it at all and stop using it until it's fixed. Computers aren't designed to perceive the "A" and check that it is an appropriate result in the light of the command. They are built so that such perceptual checking is unnecessary. That's the main reason that computers are not like organisms.

That's also why computer-intelligence types have taken so long to discover feedback. It never occurred to them that the output of the nervous system might be different from what some higher center commanded, requiring correction. It never occurred to them that a command "lift the block" might fail because the output device doesn't apply a force straight up, or doesn't apply the same amount of force every time, or because something else is pushing sideways. It never occured to them that a general command might be ambiguous when translated into more specific actions, so that something else -- SOMEHOW -- would have to pick one of the possible implementations that would in fact satisfy the command. In a top-down system, what is it that knows that a particular lower-level output will in fact satisfy the command?

>i believe none of the reflective levels we've been discussing is
>going to help us figure out what is really going on. we need to

>dump them and look for new levels (if any).

Try my proposals. They aren't about language, but about levels of experience.

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

Come on. People quite commonly use sentences as part of their language processes. You've made your point that language doesn't necessarily come in sentences. I've tried to show that there's more to language than anything made of words, but we do use conventions in our communciations, which relate directly to the way we organize words for output. The conventions don't always involve sentences, but conventions using sentences exist and are used. People do language differently in different circumstances. When you do it with writing, you usually do it with sentences, unless you're a poet or a songwriter. A lot of people speak in sentences, if they're not politicians. Is this just an affectation? I think it's one way to do language.

Best, Bill P.

Date: Fri Oct 30, 1992 5:03 am PST
Subject: cat getting a licking

[From: Bruce Nevin (Fri 921030 07:47:34)]
(From: David Goldstein Date: 10/29/92) --

I've seen cats lick themselves when they're nervous about something that they want to keep an eye on but (as I have interpreted it) without wanting to appear as if on their guard. The value of this, I imagined, was that going one's guard in an evident way might be interpreted as a hostile act by some parties that one might be nervous about. In these sorts of situations, the licking is not at all relaxed and indolent, as it often is in other situations. A bit more provocation and the cat might abruptly get to its feet, but then walk away slowly, as though unconcerned, but (as indeed with the nervous licking) with perceptible tension and a (my) sense of it being irritated. More than that and the departure from unpleasantness is less demure, graduating to full blown alarm and retreat.

Just my interpretation of cat watching.

(Being licked by its mother was presumably reassuring).

Bruce bn@bbn.com

Date: Fri Oct 30, 1992 11:33 am PST
Subject: Re: SR and "sorry"; control of meaning; social control

From Oded Maler (921030) * [From Bill Powers (921024.0830)]

* I'm not sure how you see the idea of stimulus and response entering
* into tracking a cursor. In our experiments there is always a
* disturbance, so "responding to the stimulus in a certain way" will not
* result in tracking. There's no response that corresponds to each
* target position such that the result will place the cursor at the
* target. If you consider the cursor as part of the stimulus, then it's
* neither a dependent nor an independent variable.

My personal meanings of stimulus and response are much less loaded than that of psychologists. What I meant was that in some moment the tracking system receives a reference signal that makes it respond in certain way (that is, try by variable means to achieve a certain perceptual goal in the presence of disturbances, bla-bla, etc.). Perhaps in your setting this event occurs only once when the experimenter tells the subject the rules of the game ("try to track everything that appears"). Now suppose that the figures appear and reappear on the screen, and that after some training the rules change: some shapes you have to track, but others you ought to ignore. All this was just a suggestion of how to reproduce a "sorry"-like effect in the laboratory.

* Perhaps I didn't make my question clear. When I asked what kind of
* experiment you could do to determine that the "sorry" loop is lower in

* the hierarchy than the conscious loop, I was only asking what criteria
* you would use to determine that saying "sorry" was lower in the
* hierarchy than some conscious kind of behavior, such as evaluating
* whether saying sorry is appropriate.

I think hierarchies are personal and cultural. I can observe some parts of it in, say, species from the French culture, because I'm an outsider from a culture where this reflex (and others, such as looking backwards before you close the door while entering a building) is situated in the higher conscious level (if at all..) The ideal way to determine it would use a crude hypothetical (and probably impossible) mapping of the hierarchy to the brain, and that use some lesions/drugs that are known to affect certain areas (areas in the broad sense of a generalized geometric/chemical topology). Till then I don't see any *scientific* way to speak about it.

* You must have some sort of actual behavior in mind, and use some way
* of interpreting what you observe. How else can any theoretical concept
* have meaning?

They have private meanings to me as a pertt time theoretician of my own mind.

The question of how to determine whether "something" is higher or lower leads to the general problems of whether CEVs really exist (otherwise, what is that "something" in the subject's mind that you and me are talking about?). In your experiments, you assume that the percept of, say, a cursor on the screen has some universal characteristic so that you may assume the existence of a corresponding signal somewhere in the hierarchy of every subject, as well as the possibility to refer explicitly to this signal by verbal means, while instructing the subject.

Well, I don't seem to converge to an answer. The answer should be however related to the notions of conflicts and bounded resources within levels.

--Oded

Date: Fri Oct 30, 1992 12:07 pm PST
Subject: Re: Language perception and production

A late comment on sound shifts. I was tempted to write net along the following lines earlier, but I had missed this posting. Perhaps the issues are illuminated by the kind of argument George Sacher used to use in evolutionary biology. Roughly speaking this goes as follows for the sound shift case:-

There is a tradeoff between cost of clear articulation needing precise conscious control and error rate on the part of listeners.

We can classify conversations into classes by where this tradeoff lies. Obviously a commencement address has a different optimum, from ordinary conversation, and this has a different tradeoff from being a small LRP group, way out in enemy teritory in the jungle.

Within a class there is still variation, and the usual stuff in population genetics says that rate of evolution or divergence is proportional to variance. I don't remember if this is the Hardy Weinberg Thm. or something else, but could probably look it up.

So, it pays to look at redundancy and error rate in perception of phonemes and do a Shannon kind of analysis to estimate acceptable variance.

Refs - Ref in net posting to work of Slawinski and of Ohala

Date: Sat Oct 31, 1992 7:26 pm PST
Subject: CSGnet Slow-downs

[from Gary Cziko 921131.0510 GMT]

From time to time CSGnetters may notice a slow down in CSGnet traffic. There are two possible reasons for this. The first is that no one or few people are sending messages to the network. The second is that the machine here which manages CSGnet is low on disk space due to lots of unread electronic mail. When this happens, no messages are sent out on any of the many lists managed by the machine since it is "afraid" that there will be no room to store the resulting "ricocheting" mail. Sending resumes when enough disk space frees up. Mail sent to CSGnet during these times will almost certainly be distributed when the network returns to life again. But if you want quick confirmation, you might want to wait until you receieve CSGnet mail again before posting.

Both of these cases are beyond my control and there really is nothing to do but wait. During these times I often get many messages from addicted-beyond-all-hope CSGnetters asking me nervously if "something is wrong" with the net. Instead of sending such messages, I suggest that when things slow down you just sit back and relax until it starts up again. Read that book you've been wanting to. Write that book you've been hoping to. Take your significant other out to dinner. Read a book to your kid. Enjoy the breather, since before long you'll be buried again in CSGnet discussions and you'll have to read as fast as you can just to avoid falling further behind (the maddening Red Queen Effect). --Gary