Date: Mon Feb 01, 1993 4:10 am PST
Subject: Feed-back in various levels

[From Oded Maler (930201)] Bill Powers (930129.1300)

Just to clarify, I didn't "accuse" your arm model for neglecting the environment. I was trying to give a meaning for "being too slow for.." and then observed that there are two variations on the meaning. One is "too slow as a model to explain an observed behavior" (this is Rick's emphasis I think). And the other is "two slow to be embedded in a system that achieves some 'objective' performance criterion" (this is more engineering/robotics emphasis). I tried to elaborate on the latter, and then the difference between these two questions became more clear to me.

As your comments show, it might be that in many cases if a feed-back based system cannot achieve an external objective goal (because of theoretical limits), than maybe non -feed-back based systems can't do it either. It is probably more true in the lower levels. In the higher levels the time-scales are non-uniform. Suppose one controls for "being happy" by means of "getting a tenure" - what is the time-scale here? Or taking it a bit lower: I want to go to Paris, my lower-levels deal with whatever disturbances (weather, train tickets, forgot something, etc.) as long as the disturbances are reasonable, I control for Paris in a feed-forward manner, but suddenly, near Lyon, there is a strike at the railway station, I feel hungry nervous and tired and then start to think, whether I should go to Paris at all under this conditions - this is feed-back.

--Oded

Date: Mon Feb 01, 1993 9:10 am PST
Subject: Personal mail on CSG

It is kind of confusing to receive mail like this:

>Bill, glad to send you a reprint of "Integrated Data ... Systems ... "
>It's on the way.

from someone I don't know. Then I of course realize that it isn't for me even though it is in my mailbox and addresses me as Bill, it's for Bill Powers. The result is some wasted time and irritation.

Please, if you are sending mail to Bill Powers, send it to Bill Powers. If you want to copy it to the full CSG mailing list, please identify the principal recipient as Bill Powers, not just as Bill.

I suspect that there may be list members with names like Martin and Gary
who have encountered similar problems.
--
Bill Silvert

Date: Mon Feb 01, 1993 9:17 am PST
Subject: Re: words & meanings

[Martin Taylor 930201 11:45] (Bill Powers 930129.1300)

>>Words don't have meanings so much as tendencies which coalesce
>>in the course of a dialogue (or monologue) to have some effect
>>on the listener (reader).
>
>I'm afraid that that way of putting it gives too much comfort to
>those who want to endow words with special properties. Words
>don't have EITHER meanings OR tendencies. Saying it that way
>ignores the active agent that is using the words for a purpose.
...
> Words have no tendencies at all. They are perceived
>or they are not perceived. That is all they can do. The rest is
>up to the receiving system.

Quite right. Correction accepted. That's one of the bases of Layered Protocol Theory. But the point I was trying to make was that what the perceiver does with the words is not specifiable by any precise definition. Rather, the perceiver (who may be the talker) deals with the ways the words tend to be used, in such a way that their effect is dependent on the verbal and non-verbal context. Most particularly, it depends on the perception the perceiver has of the current state of the communicative partner, including models of what the partner would be likely to do with the word.

There was no point in going into the whole of Layered Protocol theory, just to make the point that natural languages are not mathematical symbol systems. But you are right, the words themselves have no "tendencies." But most talkers of a language will tend to use any given word in much the same way, so I think it isn't very misleading to talk as if the words had meanings or tendencies toward meanings.

Martin

Date: Mon Feb 01, 1993 10:36 am PST
Subject: Re: FB 2 slow; words & meanings;devil's bib

[Martin Taylor 930201 12:20] (Bill Powers 930129.1300)

On measuring transport lag:

A little background. If a receiver is supplied with a variable signal, the amount of information that receiver can get from the signal is bounded by its prior knowledge of the signal characteristics and by the power and bandwidth of the signal. If a receiver is simply detecting which of two events occurred, according to what it receives in that signal, its accuracy is limited by the same parameters. It is possible to show that there is a relationship between the accuracy, measured by d', and the maximum information that could be obtained through any set of signals of the same energy. The relation is $d'^2 = 2C \ln 2$, where C is the amount of information in bits.

If you ask a subject to see which of two lights comes on, and to press the appropriate button, the answer will (almost) always be right if you give the subject enough time. But if you make the subject respond at a particular moment, the answer will be a pure guess if that moment is before the light comes on, or for some short time thereafter (you can set the moment pretty accurately by training the subject to press the button on the third of three quick "bip" tones). It turns out that if you plot d'^2 for the choice of the correct button as a function of when the button was pressed, the plot is a very accurate straight line with an asymptote somewhere around 210 to 230 msec (depending on the subject). If you want to talk about PCT fits to data being accurate, this one requires you to plot the line with gaps if you want to see the actual data points (Taylor, Lindsay and Forbes, Acta Psychologica, 27 (1967), 223-229).

The linearity of the function suggests that in this particular configuration, the observer was getting about 140 bits/sec over at least the next 100 msec, but that this information starts to be available to whatever mechanism affects the button choice only after about 210 to 230 msec. This isn't, of course, the real transport lag, since the action of button pressing was timed by the "bip" sequence. But it is some measure of the time before which no perceptual information can be used in the act of selecting the button

Martin

Date: Mon Feb 01, 1993 12:15 pm PST
Subject: feedback too local

[Avery Andrews 930202.0630]

In the motor control literature I've run across statements to the effect that feedback is `technically' restricted to monitoring immediate products of the effectors (such as, presumably, rpms on a driveshift). Hence, if you are trying to close your lips and compensate for one being disturbed by moving the other further, this is `technically' not feedback.

Do any of the engineers on the net recognize this as an actual doctrine from courses they took, or anywhere else? Or is it perhaps just an idea that psychologists picked up somehow?

Avery.Andrews@anu.edu.au

Date: Mon Feb 01, 1993 2:37 pm PST
Subject: Good Data, No Model

[From Rick Marken (930201.1200)] Martin Taylor (930201 12:20) --

>It turns out that if you plot d'^2 for the choice of
>the correct button as a function of when the button was pressed,
>the plot is a very accurate straight line

> If you want to talk about PCT fits to data being accurate,
>this one requires you to plot the line with gaps if you want
> to see the actual data points (Taylor, Lindsay and Forbes, Acta
>Psychologica, 27 (1967), 223-229).

I have Martin's article and I agree that the fit is very impressive. But I'm not quite sure what it means because I don't know the model that is being tested. What is the mechanism that produces the very precise relationship between delay after stimulus presentation and d'^2?

I suspect that the results have something to do with the level of the perceptual variable that is being controlled; for example, in order to be right all the time the subject must be able to perceive a relationship ("light on the right", "light on the left"). But the subject might be able to be right some of the time based on lower level information ("the more intense light" might sometimes be enough to indicate "the light on the right"). The relationship perception (which is a sure-fire indication of which "response" to make) requires more time to occur than the intensity perception (which is a less reliable indication of the appropriate "response"). So what you are seeing (I hypothesize) is a side effect of what I demonstrate in my "Hierarchical behavior of perception" demo -- the time constant for higher order perceptions is longer than that for lower order ones.

But a working model is absolutely necessary here in order to know what might be going on; this is a good example of why even high quality data is of little value (well, it is valuable inasmuch as it raises questions; but it is not at all clear what the data means) if it is not based on exploration of a working model.

Best Rick

Date: Mon Feb 01, 1993 6:22 pm PST
Subject: feedback too local

[From Rick Marken (930201.1730)] Avery Andrews (930202.0630)

>In the motor control literature I've run across statements to the effect
>that feedback is `technically' restricted to monitoring immediate
>products of the effectors (such as, presumably, rpms on a driveshift).
>Hence, if you are trying to close your lips and compensate for one being
>disturbed by moving the other further, this is `technically' not
>feedback.

Yes! This is what I was trying to get at with my post about converging and diverging "bushes" for perception and output respectively. What you have noticed, I believe, is the failure of motor control (and virtually all other) psychologists to get one of the most important points of "Behavior: The control of perception". The model in that book was presented largely to show the feasibility of the notion that ALL BEHAVIOR is controlled perception. The notion is feasible when we see that higher order perceptual variables can be controlled via control of lower order ones. I developed the spreadsheet model just to show that what Bill P. said in BCP also works when you carry out the computations; the spreadsheet model controls logical perceptual variables (a>b = true) by controlling many different lower order variables. If psychologists could get the idea that higher order variables (including variables that change over time) can be controlled then they might see the richness and depth of PCT and, maybe, pay some attention to it. The won't do that as long as they think of feedback as being restricted to the "immediate products of the effectors".

A very good observation, Avery.

By the way, the reaction time stack is coming along -- it's just VERY slow in Hypercard.

Best Rick

Date: Mon Feb 01, 1993 7:32 pm PST
Subject: Re: feedback too local, 2 output blunders

[Avery.Andrews 930202.1422]

Thinking about the `output blunder' today, it seems to me that maybe my version and Bill's really are the same thing after all, namely confusing the product of the effectors (say, some forces) with the result that is desired. Neglecting varying `disturbing' forces is one aspect of the blunder, neglecting the almost always indirect, tho sometimes relatively constant and predictable relationship between the effector product and the result is another (e.g., the connection between moving into the shade and having less heatflow into your body, if you're a thermoregulating lizard).

Perhaps `result blunder' would be a good name for it, if its not already being used.

Avery.Andrews@anu.edu.au

Date: Mon Feb 01, 1993 10:38 pm PST
Subject: Levels; feedback too slow; open-loop models

[From Bill Powers (930201.1900)] Bob Clark (930131) --

>I am leaning toward designating Seventh Order as the Order
>including perceivable variables off Personality. This would
>imply considering Eighth Order as pertaining to Character.

I'm glad you put it that way, because it brought into focus a difference between the way you're characterizing higher orders of organization and the way I'm thinking of them.

To speak of "personality" and "character" is to take an external view of someone else's organization. That is, you seem to be looking for levels that will apply to "psychological" aspects of a person, to explain the how and why of that person's behavior.

I'm taking a different viewpoint: my definitions of levels are meant to describe how the world appears from the standpoint of the person regardless of the context. When I speak of "system concepts," I'm referring not just to things like a self or a personality or a character, but to ALL system concepts. To a physicist, for example there exists something called physics, a discipline. This is, of course, a perception. The entity called physics, I have proposed, is a concept build from a set of principles and generalizations, which both provide the material within which the entity physics is perceived, and which, as goals, are specified by the goals we have for physics -- that is, for what kind of entity we want it to be.

The principles and generalizations, in turn, are built out of a set of rational, logical, reasoned mental processes that I call, generically, "programs." In a set of programs we can discern general principles; at the same time, the principles we wish to maintain in force determine what programs we will select to use.

My intention in proposing these levels of perception was to provide a framework within which we might understand all human experiences, no matter what they are about. If the subject matter is one person's experience of other individuals, then what I call "system concepts" would correspond to what you term "personality," and perhaps what I call "principles" would correspond to your "character," and my "programs" to something like "habits or "abstract skills" or "intelligence." These are ways of perceiving other people.

But these general classes of perception and control include more than our experiences of other people. As I said, they include all experiences of all kinds. To a manager, the system concept called "my company" is as much an entity as "my children." To a patriot, "my country" is a real living entity. To a sociologist, "society" is a system concept with as much reality as "self." And to a chemist, chemistry is an entity with characteristics that depend on principles that are implemented as programs, without any organisms in the picture.

So what I am most interested in are the general classes of experience, not specific contexts in which we might give them more specialized names. The concepts of "character" and "personality" are inventions, but they are examples of fundamental classes of perception shared by the educated and the uneducated alike, and constant across cultures (I sincerely hope).

>I am reminded of the "Leader-Follower" demo with the Portable
>Demonstrator. Do you remember this, Bill?

Yes, indeed, and thanks for bringing it up. It's been a long time since I mentioned it, however, and it really does belong in Gary Cziko's collection. Just to expand a little on your brief description:

The object of the demonstration is to see how long it takes people to switch roles; namely, from leader to follower. B moves a finger arbitrarily in space while A tries to keep a forefinger aligned with B's finger. This results in B tracing out some pattern in space, while A's finger lags behind it a little, always trying to catch up.

Then, on a signal from a third party, the two participants swap roles. Now A is moving a finger in arbitrary patterns while B tries to track it with a finger as closely as possible. Clearly, it is now A who creates an arbitrary pattern in space, while B's finger lags a little behind it, always trying to catch up.

The third party keeps giving the signal at variable intervals, and the participants keep swapping roles, until they are executing the swap as fast as possible. The claim that Bob and I would make is that the minimum possible time required for this swap is longer than the time taken to change any lower-level control process.

The time should be longer, for example, than the time required to correct the error when tracking a regular pattern over and over, with the disturbance being a sudden stop in the target pattern. And that time is longer than it takes to track a target that moves in random jumps to fixed positions, which is longer than the time it takes to respond to a downward push by swinging the arm rapidly downward, which is longer than it takes for a directly-disturbed arm to begin to move back toward the undisturbed position. So we would seem to have five nested and demonstrable levels of control with progressively longer reaction times, the fifth being the role-swapping and the lowest being the position reflex. A proviso is that all these tasks should be well-learned so we aren't looking at reorganization along with the control actions.

I just checked this out with Mary, and it still works. While checking it out, it occurred to us to wonder what would happen if one of the people simply changed roles without warning the other and without any external signals. With different pairs of people, the results might be different, but in our case the result was hilarious. I wons't spoil it for you by describing it.

Rick Marken, you could surely program a computer to do this -- just run a model that makes the target track the cursor, or at random intervals, switch to generating a random 2-d pattern to be followed.

Avery Andrews (930201.0930) --

>But is this really a blunder? After all, if an `output' is
>used as a reference level for a control system, you do get a regular result.

How are you going to use an output for a reference level? I think you might have meant to say "if a reference level is set for an output." But even if that's what you meant to say, it's still not right, because what we control for are perceptions, not outputs. The outputs vary as disturbances require. We perceive and control consequences of outputs mixed with consequences of disturbances, not the outputs themselves. All of our outputs are accomplished by shortening the contractile part of a muscle, a phenomenon that we neither perceive nor control. Even at the lowest level, all we perceive is a consequence of that shortening: a change in a tendon receptor's perceptual signal, or of a stretch receptor's perceptual signal. From that level, we proceed to controlling perceptions of consequences of controlling those signals, and perceived consequences of those consequences, and so on. All of those are inputs, not outputs.

If we forget to insist on that description, which as far as I know is always rigorously correct, we risk losing the whole distinction between PCT and the older way of viewing behavior. As soon as we start allowing some slack in this description, we fall into confusion, and lose this essential distinction.

>It is also my impression that people have in fact pretty much abandoned the
>`output regularity' blunder for most kinds of movements, at least.

If so, this is news to me. And how about talking about consequences of movements? Isn't it generally assumed that if one can reliably produce a series of preplanned joint angles, that motor behavior in general is pretty much explained? But just try applying that model to the limb movements that turn the steering wheel of a car, computing the motor program needed to keep the car on the road -- without feedback.

Unless you can convince me otherwise, I still believe that the output-regularity blunder is committed by all motor program and central pattern generator models. It's assumed that if somehow you can produce just the right pattern of command signals, the limb will move to a corresponding configuration in space, without feedback. And the converse is the assumption that if the arm does move to a particular spatial configuration, it must have been commanded to do so by a set of signals on which that position regularly depends. This concept requires that muscles never fatigue and that no extraneous forces ever act on the arm in unexpected or unpredictable ways.

Perhaps Schmidt has realized that motor programs don't work. How about Bizzi and Kelso and all the other biggies who are still publishing the same old stuff? When you say "people have abandoned" the motor program or pattern-generator view, which people do you mean? I think you mean another minority like us.

Gary Cziko (930131.1705) --

RE: Ballistics vs. feedback demo.

I agree with Avery that your throwing demo makes the point more directly than the coin demo does. However, it inspired me to think up another one that shows a little more of the effect you want.

We happen to keep around the house various toy trains, for purposes of grandparenthood. I picked a wooden train car about 6 inches long, weighing about a pound, with wooden wheels and a convenient hook at each end. To each hook I fastened a string of 3 rubber bands, fairly weak. I set the car on a table, on its wheels. Then, using both hands, I stretched out the rubber bands so the car came to a balance point between my hands.

There are now two ways to move the car. (1) Move both ends of the rubber bands by a fixed distance to one side and let the car end up where it will; (2) watch the car and move your hands (keeping tension between them) so as to bring the car to a fixed position.

If you try to move the car as fast as possible by method (1), you can make two marks on the table and move your hands to the marks as rapidly as possible. The car will be accelerated in the direction your hands move, reaching maximum velocity just as the tensions in the two rubber bands are equal. It will then proceed past the midpoint until its velocity is reduced to zero by the growing tension in the trailing rubber band. It will then accelerate back the other way, and so on in diminishing oscillations to an end-point.

Using method (2), you mark the final position of the car resulting from method (1), then reset the car to its original position. Now you watch the distance between the car and the mark, and move your hands in parallel, maintaining tension between them, to bring the car to the mark. It will move to the mark and stop there with no oscillations. With practice you can make it do this far more rapidly than you can get the car to the mark the other way.

This would be even more dramatic if the rubber bands were very weak and the wheel bearings good. You would have time to accelerate the car toward the final position by moving your hands far to one side, to get a strong acceleration, and then far in the other direction to slow the car to a stop, your hands returning to the correct final position automatically.

The only way to make method (1) work almost as well as method (2) would be to generate an arm-movement waveform just right to produce a high initial acceleration, and then at just the right time, a high final deceleration. In other words, to provide a central pattern generator of high precision that produces the same arm-positions that the control system generates without preprogramming.

This is why people have been driven to proposing motor programs instead of systems that just issue a "position" command. The motor programs are supposed to compensate for the dynamics of the controlled variable, as well as the kinematics of the jointed limbs. Once you start down this trail, still thinking of commanding output, you are driven step by logical step until you fall into the hole. Your basic premise leads you to propose a pattern generator of incredible precision, and a program of equal precision that bases its command outputs on unobtainable data of just as great precision -- and it requires you to ignore all long-term disturbances.

Method (2) is not only simpler and faster, but it can work
indefinitely (no cumulative computation errors) and it can
achieve good final precision using low-precision output
effectors, even in the presence of environmental nonlinearities
and disturbances.

Avery Andrews again (930201.1059) --

>So, the jab might be driven thru a kinesthetic control system,
>but it also might be that people temporarily deafferent
>themselves at low levels (remember the Bizzi monkey arms), and
>the fact that it requires so much practice to acquire these
>moves makes anticipation, etc. much more plausible as a
>mechanism than it normally is (as Tom Bourbon just pointed out).

"Anticipation," I take it, is perceiving the approach of something that is about to disturb the arm and preparing an output that will occur just as the disturbance occurs, cancelling its effects.

Of course for this to work, you must practice until you can estimate the exact effect that the disturbance is going to have on the arm, and you must calibrate the output command signal so that its effects will cause the muscles to tense along a curve that is opposed to the disturbance's effect and matches the rate of rise of the effect as the disturbing object comes into contact with the skin, imparts an impulse to the arm, and bounces away.

I agree that this can be done, and is done. It doesn't work very well, but it does reduce the initial effect of the disturbance somewhat. Ask someone to suddenly drop a book onto your outstretched hand from a height of about 2 inches. You'll see a disturbed control system. Then ask them to drop it from a height of a foot. You'll probably see your hand RISE to catch it, and then lower it back to the initial position. This looks to me like a higher-level system, which acts more slowly but more intelligently, readjusting the lower-level (alpha-gamma) reference signal in a sort of compromise: creating an upward position error to minimize the downward position error after the disturbance.

I still don't buy the idea of people deafferenting themselves. It seems to me that that process would take longer than the control process it's supposed to speed up. And there is absolutely no evidence that I know of that the basic spinal loops are ever turned off save by relaxing completely (zero reference signals to everything).

>I think it's important not to spend too much energy on
>peripheral matters where PCT expectations might not be borne
>out -- if you insist that the jab is done via kinesthesis, and
>somebody proves that it isn't, then you've lost a lot of credit.

The whole point of making models is to make them testable, so they CAN be disproven if nature so chooses. I'm not out for the kind of credit you're talking about (speaking so cautiously and generally that you don't risk being wrong). I think that sort of CYA approach is one of the main things wrong with the life sciences, and the main reason that they don't have any good models.

Gary Cziko (930201.0402 GMT) --

Avery's explanation of inverse dynamics is right. A footnote.

In the few instances of inverse dynamics that I've seen, there's an assumption of a certain path with initial position, velocity profiles, and final position given. From the characteristics of this desired outcome, and taking into account the geometry (kinematics) and inertial properties (dynamics) of the jointed arm, the model calculates the waveforms of the torques that must be applied to the joints to produce that outcome. If unique inverses are found, the problem is assumed to be solved. But there is one more step required.

That step is to take the computed waveforms of torques, apply them to a model of a jointed arm using the FORWARD kinematics and FORWARD dynamics, and compute what the arm actually does.

From a purely mathematical standpoint this second step is superfluous; if you found the inverse of the desired path, then going the other way can only result in that same desired path. But from the modeling standpoint, the second step is critical, because if you actually do both the backward and the forward computations using a computer, as the nervous system must actually do them, you will NOT end up with the desired path, but only with something more or less similar. The hundreds of arithmetic calculations involved will produce errors both in computing the required pattern of torques, and in converting that pattern back into the resulting movement profile.

Suppose we run the computed arm position back and forth between two positions, say 1000 times in a row. Now no amount of feasible precision of computation would prevent a drastic drift of the path away from tracing and retracing of the desired profile (a closed-loop control system could move back and forth between the points indefinitely with no loss of accuracy).

But even that's not the worst of it. If this model is supposed to represent a real arm, then we must suppose that the precision and repeatability of the torques reflects that of the nervous system and muscles. If we reduce the accuracy of just the output apparatus to a realistic level, leaving the central program computations with infinite accuracy, the forward step in testing the model will fail cruelly. The reason is that torques create velocities and velocities create positions: two integrations, each one hypersensitive to initial conditions. We have not only rounding errors accumulating in both directions of the computations, but a double integral of the rounding errors, which magnifies itself drastically over time.

If those who propose computation of inverse arm properties had used their model to compute torques, and then actually used the computed torques to compute the predicted behavior of the arm -- in other words, if they had actually tested their model -- they would have realized the enormous problem of precision that is entailed by their basic design.

As long as you only do the exact analytical mathematics, you will never discover this problem.

Oded Maler (930201) --

>I was trying to give a meaning for "being too slow for.." and
>then observed that there are two variations on the meaning. One
>is "too slow as a model to explain an observed behavior" (this
>is Rick's emphasis I think). And the other is "two slow to be
>embedded in a system that achieves some 'objective' performance
>criterion" (this is more engineering/robotics emphasis).

OK, I see. My approach slices the pie differently. The claim generally made in the literature is that specific behaviors can't be done with feedback as fast as they are observed to be done, so they must be done without feedback. There's an assumption that a feedback process somehow adds complexity to the production of the action, and that this complexity prevents fast action.

Just from inspection of the canonical feedback diagram that we use, it's clear that a change in the reference signal is transmitted immediately, through a single synapse with the comparator, to the output effector. Therefore the most time that could be saved by eliminating the feedback would be the synaptic delay in the comparator, on the order of one millisecond. In a robot the savings would be even less -- a few tens of nanoseconds at the most. This small time savings would be traded for a great loss of precision.

>As your comments show, it might be that in many cases if a
>feed-back based system cannot achieve an external objective
>goal (because of theoretical limits), than maybe non -feed-back
>based systems can't do it either.

Yes, and I think this conclusion is general. If the system isn't monitoring the approach of something to an external goal- condition, how does it know when to stop acting? And if it is monitoring the external condition, then it's really controlling a perception, isn't it? Knowledge of the current external state of affairs IS feedback.

>In the higher levels the time-scales are non-uniform. Suppose
>one controls for "being happy" by means of "getting a tenure" -
>what is the time-scale here?

Quite long. If one cannot perceive one's state of happiness, and can't do anything to affect the state of tenure of which happiness is a function, then the time scale is infinite. You can't control for tenure if you can't perceive how close you are to getting it, and you can't use achievement of tenure to control for happiness if you can't perceive how happy you actually are, compared to how happy you intend to be. Lots of people discover that when they get tenure, their state of happiness actually decreases: i.e., the supposed dependence on tenure was imaginary, and they were already happier than they supposed they would be after getting tenure. You have to pay attention to your perceptions if you want good control of anything.

>I want to go to Paris, my lower-levels deal with whatever
>disturbances (weather, train tickets, forgot something, etc.)
>as long as the disturbances are reasonable, I control for Paris
>in a feed-forward manner ...

No, you control by controlling what you can control: being in process of going to Paris. This goal you can achieve immediately, by making the first move of your hand toward the telephone to call the travel agent. If a friend were to ask you what you're doing as you talk to the agent, start throwing clothes into a suitcase, check your wallet for money and credit cards, call a friend to feed your cat, what would you say? "I'm going to Paris." All these lower-level details are perceived as a process, and the process matches your reference level for it as long as it's going right.

What you don't do is say "I'm in Paris" and expect simply to be there. That isn't how control works. You have no means of just suddenly being in that place you imagine. You have to put Paris into a process as the end point, and construct a path, and get on the path, and stay on the path, resisting disturbances that might take you off it. As long as you're on the path, you're satisfied that you're achieving the goal of going to Paris. When that process has been guided through all its parts to the end, you can then examine the next attainable goal: get some good food. That entails constructing another process-goal, which you can immediately achieve AS A PROCESS.

Human beings aren't necessarily very good controllers at the higher levels. They do in fact set unattainable goals, not realizing that they've done so. They say "I want to be a millionaire." But that goal can't achieved by a normal control process. People with just that goal never become millionaires by their own efforts. They play the lottery.

If someone wants to be a millionaire, the only way to do it is to put the state of having a lot of money at the end of a path that entails doing and learning to do all the things that accumulate a lot of money fast. Then you have to get on that path, start learning and doing. What happens then, of course, is that the original process, for some, becomes interesting in its own right; it brings up all sorts of problems that have to be solved to stay on the path; it demands full-time attention, because generally other people resist giving up their money in large quantities. By the time the end-point is reached and the bank account reads $1E6, the person has become immersed in the process and has forgotten its original purpose. This accumulation process has become a parasite with its own autonomous existence. Fortunately this disorder affects few people. Unfortunately, it makes them very powerful.

How did I get onto that? Probably from looking at my checkbook balance.

Martin Taylor (direct post)

We'll get more info to you in late March or early April. On the pledges, you can send yours to me personally and trust me to apply it at the right time in the appropriate manner (for everyone else: a problem with cashing foreign checks at the meeting, and how to pledge money to help subsidize more students. $200 accumulated so far, almost enough for one more student).

(930201) --

Your experimental results sound extremely impressive. I am impressed. That's what I call good data.

How are you estimating C in computing d^2?

Avery Andrews (930202.0630)--

>In the motor control literature I've run across statements to
>the effect that feedback is `technically' restricted to
>monitoring immediate products of the effectors (such as,
>presumably, rpms on a driveshift). Hence, if you are trying to
>close your lips and compensate for one being disturbed by
>moving the other further, this is `technically' not feedback.

>Do any of the engineers on the net recognize this as an actual
>doctrine from courses they took, or anywhere else? Or is it
>perhaps just an idea that psychologists picked up somehow?

As far as I know they didn't pick it up anywhere; they invented it. Some people can't tell the difference between a plausible inference and knowledge.

Actually when it comes to defining terms there are no rules; you can call moving one thing relative to another to keep the relationship constant "indirect multiple-mode reflexive adjustication" if you want. When you lay out a model that can do this and write the equations, they will have an uncanny resemblance to the model and the equations used for "technical" feedback control systems, and the general behavior will be indistinguishable from that of any other control system. But if you don't want to call this effect feedback, you are perfectly right to use some other term, unless you want people to think that you know something about the subject.

Best to all, Bill P.

Date: Tue Feb 02, 1993 1:14 am PST
Subject: technically feeedback

[Avery Andrews 930202.2003] (Bill Powers (930201.1900))

>As far as I know they didn't pick it up anywhere; they invented
>it. Some people can't tell the difference between a plausible
>inference and knowledge.

I certainly haven't been able to find any warrant for it in the control theory books I've looked at - I conjecture that they picked it up because it is more or less true of many of the actual examples discussed in the books, in spite of its nonessential and fundamentally incoherent nature.

Avery.Andrews@anuj.edu.au

Date: Tue Feb 02, 1993 7:31 am PST

Subject: Asking the devils; transient control

From Greg Williams (930202) Bill Powers (930201.1900)

>If so, this is news to me. And how about talking about
>consequences of movements? Isn't it generally assumed that if one
>can reliably produce a series of preplanned joint angles, that
>motor behavior in general is pretty much explained?
>...
>Unless you can convince me otherwise, I still believe that the
>output-regularity blunder is committed by all motor program and
>central pattern generator models.
>...
>Perhaps Schmidt has realized that motor programs don't work. How
>about Bizzi and Kelso and all the other biggies who are still
>publishing the same old stuff? When you say "people have
>abandoned" the motor program or pattern-generator view, which
>people do you mean? I think you mean another minority like us.

Bizzi and Mussa-Ivaldi, "Muscle Properties and the Control of Arm Movement," in Daniel N. Osherson, Stephen M. Kosslyn, and John M. Hollerbach, eds., AN INVITATION TO COGNITIVE SCIENCE, VOLUME 2, VISUAL COGNITION AND ACTION, MIT Press, 1990:

234 - "According to the equilibrium trajectory hypothesis... the multijoint arm trajectory is achieved by gradually shifting the arm equilibrium between the initial and final positions. In this control scheme, the hand tracks its equilibrium point, and torque is not an explicitly computed variable."

239 - "This work indicated a planning strategy whereby the motor controller may avoid complex computational problems such as the solution of inverse dynamics [there, Gary!]. According to the equilibrium-trajectory hypothesis, the muscle's springlike properties are responsible for generating the necessary joint torques, thus implicitly providing an approximated solution to the inverse dynamics problem. As the approximation becomes inadequate at higher speeds of acceleration, the stiffness can be increased and the equilibrium trajectory can be modified on the basis of the difference between the actual and the planned path. The task of the CNS is then to transform the planned trajectory into a different sequence of equilibrium positions and stiffnesses."

I think that endless difficulties can arise by looking at the literature, attempting to figure out what nonPCTers believe, and then critiquing those beliefs (as perceived by PCTers) without further input from those being critiqued. This is analogous to how some researchers go astray about PCT ideas -- they don't communicate very well with the source of those ideas. At the very least, it would be wise to ASK the "devils" a few questions (I assume they all have e- mail or at least postal addresses) -- like "a motor controller can be what?" -- to make sure that the words being put into their mouths accord with their current beliefs as THEY perceive them.

>No, you control by controlling what you can control: being in
>process of going to Paris.

In abandoning strict end-point control, beware the infinite regress: controlling for "being in the process of" "being in the process of" "being in the process of"... "being in the process of" going to Paris! How do you decide where to draw the line? I think it is more parsimonious to hold to end-point control, with transient conditions almost always present for the higher-level (slower time-scale) processes. Why should controlling for being in Paris be more complicated (by "controlling for being in the process of going") than controlling for putting one's finger at a particular place (modeled in the arm program by end-point control only)? It takes a (short) time to achieve satisfaction of the reference condition in pointing; it takes a (long) time to achieve the reference condition for being-in-Paris. I see no need for postulating controlling for "being in the process of," except to preserve your claim that, in general, organisms always show good control. (Actually, one might still make this claim, even while recognizing that transient, "uncompleted" control is a commonplace.

We're headed for the hospital to see how the control transient continues for my son Cambron's broken forearm. The doctor has been controlling since Sunday afternoon for seeing full wrist function (supination) within a few months. Surgery is today.

As ever, Greg

Date: Tue Feb 02, 1993 8:25 am PST
Subject: Re: technically feedback

[From Bill Powers (930202.0830)]

Avery Andrews (930202.2003) --

The idea that a relationship between two variables can't "technically" provide feedback probably arose from the fact that there is no sensory organ that can detect the state of the relationship. Higher levels of control, however, can exist. One sensory system can detect the state of A and another the state of B. The state of B can be under direct feedback control, while A can be sensed but is an independent variable, not controlled. However, a higher system can receive the sensory report on A together with the sensory report on B, and combine the two signals to generate a new perceptual signal, B-A. The higher control system has a reference signal specifying the intended magnitude of this difference. If the difference does not match the reference difference, the error signal adjusts the reference signal for B, the element of the perception that is under direct control to make the difference greater or smaller as required. Thus the perception (B-A) is maintained at its reference level. The only structural difference between the higher system and the lower one is that the higher system controls a derived feedback perception while the lower one, B, controls a signal coming out of a single specialized sensor.

The two-level system will oppose disturbances of the controlled element, B, through the action of the lower control system. Whatever error remains in the difference (B-A) due to imperfect control at the lower level will result in an adjustment of the lower reference signal in the direction that increases the lower error a little, calling forth a larger effort and eliminating (or further reducing) the error in (B-A) at the higher level. If the uncontrolled input, A, is disturbed, it will simply change accordingly, as it is not controlled individually. This will produce an error in (B-A), and the error signal will change the reference signal for the B control system to move B in the direction that will restore (B-A) to its reference level.

So: the answer to the problem you raised lies in the idea of a hierarchy of control.

Best, Bill P.

Date: Tue Feb 02, 1993 9:13 am PST
Subject: Effector-result relationship

[From Bill Powers (930202.0900)] Avery Andrews (930202.1422) --

>Neglecting varying `disturbing' forces is one aspect of the
>blunder, neglecting the almost always indirect, tho sometimes
>relatively constant and predictable relationship between the
>effector product and the result is another (e.g., the
>connection between moving into the shade and having less
>heatflow into your body, if you're a thermoregulating lizard).

The example is also an output blunder. The problem here is using a qualitative example in talking about a quantitative system. The lizard needs to keep its body temperature at some level. To be sure, moving into the shade results in less heatflow into the body. But "shade" and "heatflow" and "less" are qualitative terms, concealing the fact that the efficacy of shade is highly variable, as is the heatflow that is being "reduced." Fortunately for lizards, they do not need to regulate their body temperature as closely as mammals do. They can get by with a poor thermal control system. But it is still a quantitative system; lizards must vary the time they spend in the shade to compensate for the varying opacity of the shading object, the angle of the sun and the cloud cover, and the unpredictable amount of reflection from nearby surfaces.

The output blunder is often concealed in this way. You might say that all a person needs to do to control the position of a suitcase is to pick it up, so emitting a "picking up" output would suffice, no feedback necessary. However, the lifting force required to pick up a suitcase depends on what's inside it. It's true that a person can just pick up the suitcase without knowing what's in it. But doing that requires a position control system that can vary the lifting forces exactly as required to get the suitcase off the floor without flinging it into the air. What seems to be a simple stereotyped response when described qualitatively turns out to require precise quantitative control.

There is, in fact, hardly any kind of "relatively constant and predictable relationship between the effector product and the result." That only seems to be the case when you describe the situation qualitatively. If you pay attention to the actual quantitative physical details in just about any action, you will find that the predictability of the results would be a miracle if it weren't for feedback control.

Best, Bill P.

Date: Tue Feb 02, 1993 10:09 am PST
Subject: Re: Good Data, No Model

[Martin Taylor 930202 12:20] (Rick Marken 930201.1200)

>I have Martin's article and I agree that the fit is very impressive.
>But I'm not quite sure what it means because I don't know the model
>that is being tested. What is the mechanism that produces the
>very precise relationship between delay after stimulus presentation
>and d'^2?
>...
>... even high quality
>data is of little value (well, it is valuable inasmuch as it raises
>questions; but it is not at all clear what the data means) if it is
>not based on exploration of a working model.

Good data constrains what models might account for it, whether the data was obtained as a consequence of exploring a particular model or not. It was good astronomical observation that allowed Kepler to determine that elliptical orbits fitted as well as, if not better than, Ptolemaic epicycles, and that, in turn, buttressed (if it did not help suggest) Newton's gravitational law.

I agree that it is not clear what the data means in the absence of a model, but I would look at it from a viewpoint a little different from Rick's. I think ALL data means something ONLY in context of a model, where I take "model" in a much wider and vaguer sense than its usual definition on CSG-L. The meaningfulness of the data (I mean this in a somewhat quantitative sense) depends on the degree to which it constrains possible models. Good data is potentially more meaningful than statistically lousy data, but only if it is used in conjunction with a model that makes tight predictions: the kind of "model" that PCT people consider scientifically valuable.

No, I don't have a model of what goes on during a button choice. But whatever that model might be, it must incorporate the result that the information relevant to the choice becomes available at a linear rate after some time delay. I didn't say that this time delay is a neural transport lag, or anything else specific. But whatever model is proposed to describe what goes on in the experiment, it must provide this delay followed by a linear increase. This goes for a model to be invented in 2153, just as much as one that might have been invented before the experiment.

What is the mechanism? I'd like to know.

Martin

Date: Tue Feb 02, 1993 11:00 am PST
Subject: Re: Levels; feedback too slow; open-loop models

[Martin Taylor 930202 13:15] (Bill Powers

>Your experimental results sound extremely impressive. I am
>impressed. That's what I call good data.
>
>How are you estimating C in computing d^2?

The actual experiment wasn't mine, so I can't claim credit for the good data. It was by Jan Schouten, the founder if the Institute for Perception Research in Eindhoven. He reported them at a meeting I was at, and I used them to test the idea that people acquire perceptual information at a determinable rate.

The measured quantity is d'^2; C is derived. But the reason for plotting d'^2 instead of the more normal d' is that I had previously shown the relationship between d'^2 and C. If people gain information at a linear rate, then the plot of the measured d'^2 should be linear. And in Schouten's data it was.

Martin

Date: Tue Feb 02, 1993 12:48 pm PST
Subject: Re: Effector-result relationship

[Martin Taylor 930202 13:45] (Bill Powers 930202.0900)

> You might say
>that all a person needs to do to control the position of a
>suitcase is to pick it up, so emitting a "picking up" output
>would suffice, no feedback necessary. However, the lifting force
>required to pick up a suitcase depends on what's inside it. It's
>true that a person can just pick up the suitcase without knowing
>what's in it. But doing that requires a position control system
>that can vary the lifting forces exactly as required to get the
>suitcase off the floor without flinging it into the air. What
>seems to be a simple stereotyped response when described
>qualitatively turns out to require precise quantitative control.

But "flinging it into the air" is a common result of picking up something you thought to be heavy, but isn't. Likewise, failure of the initial lift to move something off the ground is quite common when something is heavier than expected. In both cases, there is quite a long delay before the person gets the lift under control, and in the too-heavy case, it often amounts to a total restart. Doesn't this seem to suggest that initially one does generate "a simple stereotyped response" and only after its end-point failure go into a control mode at the level of the lifting perception/action?

This ties in with the "feedback too slow" discussion, I think. Under normal condition, during a lift, feedback is obviously not too slow. But it seems to take a lot longer when the initial expectation of required force was very wrong. I'm guessing that the normal effect is that one applies a stereotyped force with the anticipation that feedback control will take over at a low level, but a higher-level control has to come into play once one has either flung the suitcase or failed to get it to move. The force goes to zero before being brought to the required level in the too-heavy case. I don't know what it does in the flinging case (maybe you have to duck, so it doesn't matter).

Martin

Date: Tue Feb 02, 1993 1:29 pm PST
Subject: Re: Good Data/No Model

[From Rick Marken (930202.1100)] Martin Taylor (930202 12:20) --

>No, I don't have a model of what goes on during a button choice.

Then why did you decide to do this experiment and collect this data? I suspect that there was a model lurking in the background; it was the ol' cause-effect model. I would venture to guess that the model was something like:

stimulus information --> processing --> output response

The linear relationship between response delay (the independent variable) and d'^2 (the dependent variable) presumably reveals something about the nature of the processing stage. The response delay was probably thought of as something that affects the amount of stimulus information or processing time available for producing the output. Is this about right?

>But whatever that model might be, it must incorporate the result that the
>information relevant to the choice becomes available at a linear rate after
>some time delay.

Not quite, I think. The model just must behave in such a way that the d'^2 measures obtained from the model (just as they are from the subject) are linearly related to the response delay measures (again obtained from the model as they are from the subject).

The problem with claiming that there is no model underlying experimental results is that it makes the experimental results themselves seem very important. But experimental results are only important (and meaning- ful) in terms of the underlying model about which they presumably provide evidence. The results per se are not particularly important. For example, it is not particularly important that there is a virtually perfect negative linear relationship between handle and disturbance variations in a compensatory tracking task. This result is only interesting because it is the kind of behavior that would be expected from a perceptual control system IN THE CIRCUMSTANCES OF THE EXPERIMENT; if you change the circumstances (change the functional relationship between handle and cursor, for example) you get a whole new relationship between handle and disturbace; but the relationship you get is still the same as the one produced by a perceptual control system in the same circumstances.

All observations (even Kepler's) are made in order to test a model; it's the model that's important, not the observations per se. I think this is an important point because it has really plagued progress in psychology. In psychology, it's the results of the research that matter, not whether or not these results confirm the underlying model being tested (for obvious reasons -- almost every experimental finding [by PCT standards] is a resounding defeat for the cause - effect model on which it is [silently] based). So we see a bunch of disconnected "facts" but no coherent picture of human nature; indeed, each "fact" seems to become the basis for it's own little "mini model" of behavior. The observation that reinforcement increases the probability of a response is a fact in its own right -- the underlying "model" being that reinforce- ment "strengthens" behavior. The observation that intermittant reinforce- ment produces behavior that takes longer to "extinguish" is another fact in itself -- with it's own little models to explain it -- even though it seems obviously inconsistent with the "model" assumed to explain the finding that reinforcement strengthens behavior.

The point of the PCT demos and experiments is that all these different little "findings" can be shown to reflect (Very PRECISELY) the same underlying process -- CONTROL OF PERCEPTION. The findings themselves are important only insofar as they test the underlying model; it's the model that's important.

So what is the finding of a PRECISE linear relationship between response delay and d'^2 other than one of the many random observations we can make about human behavior? What does this finding tell us about the organization of the system that produces it?

Best Rick

Date: Tue Feb 02, 1993 2:03 pm PST
Subject: Misc for Rick and Tom

Hello everyone,
Well, I am still waiting for Illinois Bell to give me the option to disconnect call waiting so I don't have to write messages to the net in fear that what I am writing will be destroyed--when I get it I'll get back on the net. In the meantime I have 2 or 3 questions.

First, for Rick,

A number of months ago I brought up the "error control" issue. I started with the premise that at the level of individual ECS's and at the organism level, it is Error that is controlled primarily with perceptions being controlled as a result.

Then you convinced me that at the individual ECS level, it is indeed perception that is controlled, but I still contended that at the organism level, it is Error.

You replied that the error of which I speak is a perceptual signal for another system, and hence it is still perception that is controlled. This is where the discussion stopped (but only because I was too busy to reply). Now I am curious, would you not agree that its a game of semantics once you say that this error that is controlled is a perceptual signal for another system? Of course, I have to agree that it becomes a perceptual signal--how else could it be controlled? But this is a given--my point is that the answer to the question, "What does the ORGANISM control?" is "error."

Basically what am asking is whether you agree that this is semantic quibble since I believe we have the same model in our heads. I understand that error is controlled Because it IS represented as a perceptual signal....but it's still error nonetheless.

Tom, thank you for sending me your manuscript on ERP's and PCT neuropsychology methodology. I read it a long time ago--sorry I haven't responded earlier. I understand the basic idea behind your answer to how a PCTer should do neuroscience research is to have subjects do self-initiated tasks, rather than reactive tasks. Is this correct or is there a more important point?

I agree with this approach and have been trying to understand what might be amiss in prominent articles that use a more reactive approach. Still, I do not think that a reactive approach cannot inform us. Furthermore, it does not seem to me that present neuroscience is completely immersed in reactive methodologies; there are a number of studies where the subjects' behaviors are self-initiated. And there seems to be much documented concerning disorders in which individuals cannot perform tasks which do not satisfy reducing "naturally derived" error ("Pretend like you are brushing your teeth" vs. Really brushing your teeth). These are only descriptive, of course, but they certainly seem to inform a PCT perspective.

Please keep me up to date with what you are doing in these matters and if you come across a study where the implications of the lack of reference-signal/volitional/self-initiated components are clear, I would like to know about it, for it is still difficult to articulate what is lost or in what sense misinterpretation occurs when a reactive study is performed.

On another note, this summer you spoke about study in which subjects fixate on one object and then another with a flashing light residing in between the objects. You said that after a while, the subect will have the experience of the flashing light moving to another part of space. Do you remember? Do you have the reference for this article or if not, a description of the results?

Mark Olson "It is impossible to do only one thing."

Date: Tue Feb 02, 1993 2:14 pm PST
Subject: Re: Effector-result relationship

[From Rick Marken (930202.1200)] Martin Taylor (930202 13:45) --

>But "flinging it into the air" is a common result of picking up something
>you thought to be heavy, but isn't.

I have NEVER had this experience with a suitcase or any other lighter than expected object. But thanks for the tip; I'll watch out if I'm ever checking in a very large, very light suitcase in Canada.

>Likewise, failure of the initial lift to move something off the ground
>is quite common when something is heavier than expected.

This only seems to happen when the suitcase is so heavy that some extra purchase (like a dollie) is needed to lift the thing at all.

> I'm guessing that the normal effect is that one applies a stereotyped force

How can you apply a "stereotyped force" without perceiving the force being applied. If a person were able to simply generate a sterotyped "output" -- like 200 impulses per second -- the actual force resulting from this output would not be sterotyped; in fact, it would be quite variable. In order to apply a sterotyped (about the same all the time) force, it must be the perecived force that is "applied", not the output that produces the force.

There's just no getting around it; behavior (any controlled result of neural activity) IS perception.

Best Rick

Date: Tue Feb 02, 1993 2:52 pm PST
Subject: Bizzi's model; end-point control

[From Bill Powers (930202.1200)] Greg Williams (930202) --

> 234 - "According to the equilibrium trajectory hypothesis...
>the multijoint arm trajectory is achieved by gradually shifting
>the arm equilibrium between the initial and final positions. In
>this control scheme, the hand tracks its equilibrium point, and
>torque is not an explicitly computed variable."

Here's how I interpret what I know of Bizzi's approach:

Bizzi's idea is that the arm is a mass balanced between two springs. Signals to the muscles shorten and lengthen opposing springs to change the equilibrium position. The arm is brought to the equilibrium position by the unbalanced forces. This is the same model proposed by Kelso and others of the mass-spring persuasion.

It is also exactly the muscle model we use in the arm, and we can't complain about it. However, Bizzi's model uses no feedback as far as I know, so it assumes that the observed or apparent spring constants are simply those of the muscles. If the motion of the arm appears critically damped, the damping is assumed to be that of the muscle. If the stiffness of the apparent spring changes, that is due to a change in stiffness of the muscles (due to running the opposing muscles up and down their nonlinear force-displacement curves, presumably).

My impression is that the actual spring constant and damping of the muscles is far less than what is observed with an intact system. If this is true, it would mean that Bizzi is really observing the properties of a position control system, and mistaking them for the properties of the balanced muscle pairs.

Clearly it is possible to substitute Bizzi's model (with suitable parameters for spring constant and damping) for a control-system model in which the same constants are determined by feedback sensitivities. So from the point where the command signals enter the spinal cord to the final positions of the joints, the behavior of the two models should be indistinguishable except for details of the action. Only by determining the actual muscle spring constants and force-velocity damping could we demonstrate that Bizzi's constants can't be those of the muscle, and must be set by parameters of a control system. Could this be done using data from the literature?

Having eliminated the need for computing inverse dynamics, Bizzi then goes on with the rest of the model: "The task of the CNS is then to transform the planned trajectory into a different sequence of equilibrium positions and stiffnesses."

So inverse kinematics still figures into the model. The central computer must pick a desired path, compute the sequence of joint angles necessary to produce that path, and emit the appropriate waveforms of signals to the joint-angle positioning systems. In the control-system model the inverse kinematic computation is not carried out; the same waveforms are generated using higher-level control systems, both kinesthetic and visual.

In fact Bizzi has not eliminated the need for computing inverse dynamics. The reason is that without feedback in his model, the actual parameters of the mass-spring response (based on real muscle properties) would be such as to produce inertial overshoots and oscillations for any but the slowest changes in the command signals. In order to generate a fast movement in a model with the correct muscle parameters, the central pattern generator would have to output a waveform that not only takes the kinematics into account, but also adds accelerations and decelerations as functions of time, to prevent the oscillations from appearing. Others have done this, recognizing the problem with dynamics. Bizzi gets away with his model because he is unawaredly taking advantage of the effects of the control systems, which in fact do away with the need for computing inverse dynamics. If he constructed his model using real muscle information -- obtained without feedback effects -- he would find that he can't get away with ignoring the dynamics.

>At the very least, it would be wise to ASK the "devils" a few
>questions (I assume they all have e- mail or at least postal
>addresses) -- like "a motor controller can be what?" -- to make
>sure that the words being put into their mouths accord with
>their current beliefs as THEY perceive them.

A month or so ago I sent Bizzi a letter containing the rejected Science article, the writeup of the Little Man program, and a disk with the program and source code on it. I have not even received an acknowledgement of receipt.

>In abandoning strict end-point control, beware the infinite
>regress: controlling for "being in the process of" "being in
>the process of" "being in the process of"... "being in the
>process of" going to Paris! How do you decide where to draw the line?

That's just playing with words. "Going to Paris" consists of setting and accomplishing numerous subgoals, each of which requires control to be realized in a variable environment. Some of those might also be processes, but most will not be (like reaching for the telephone).

> I think
>it is more parsimonious to hold to end-point control, with
>transient conditions almost always present for the higher-level
>(slower time-scale) processes. Why should controlling for being
>in Paris be more complicated (by "controlling for being in the
>process of going") than controlling for putting one's finger at
>a particular place (modeled in the arm program by end-point
>control only)?

Trying to characterize the means of getting to Paris as end-point control like putting one' finger on a target would be pretty naive: the output that acts directly to reduce the error would take you through the walls of the room and across field, stream, and mountain in a straight line. Moving directly toward the target works fine for touching targets when there's nothing between finger and target. It doesn't work for more complex behavior that entails strategies, like taking the train to Lyon first, in the wrong direction, to catch the Gran Vitesse to Paris.

>I see no need for postulating controlling for "being in the
>process of," except to preserve your claim that, in general,
>organisms always show good control. (Actually, one might still
>make this claim, even while recognizing that transient,
>"uncompleted" control is a commonplace.

One need for it is to account for the fact that when people are carrying out some rather complex process, they describe what they are doing in terms of the whole process, not in terms of the act they happen to be performing at that instant. The hitchhiker says "I'm hitchhiking to Chicago," not "I'm reading a menu in a Howard Johnson's", even though at the moment the latter is the truth.

But then you've never bought into the idea of levels of control, so you're sort of stuck with end-point control, aren't you?.

Hope Cam arrives at a suitable end-point. That sounds like a very nasty accident.

Best, Bill P.

Date: Tue Feb 02, 1993 4:08 pm PST
Subject: bizzi; asking the devils

[Avery Andrews 930203.1015] Bill Powers (930202.1200)

My understanding of Bizzi is that he's aware of the possibility that feedback is involved, but thinks that he's shown that its contribution for head-orientation movements is rather modest - about 10-30%. At this point I can't assess the merits of the claim. Intuitively, the mass-spring view seems unpromising for manipulating moderate-weight objects in a gravitational field. Stark's mob have a paper in JMB addressing these issues too - I don't have the exact reference right now, but random perusal through that journal is a Good Thing, I think.

Asking the devils what they mean is not going to work, since (a) they won't answer (b) it will just reinforce the image of PCT-ers as an annoying group of people who niggle about arcanities of formulation. The current literature is full of assertions as to you need reafferent feedback to deal with variable circumstances in the environment, so people think that this blunder is under control, even if they overestimate the extent to which it actually is. (`There can be no doubt that most movements are influenced by sensory input, despite some opinions to the contrary (e.g., Jones 1984)' (Cruse, Dean, Heuer and Schmidt 1990 `Utilization of Sensory Information for Motor Control', in Neumann & Prinz (eds) _Relationships between Perception and Action_.)

It's very hard to tell people something that they already think they know, so I think it would be best to leave this theme alone, and concentrate on more novel ones.

I would also not expect Arm to make that much of an impression on people until there's a lot more documentation and maybe sub-demos (of single joints, for example) that bring out (a) what is stupid about the prevailing approaches (b) what is different and good about Arm (no stored trajectories, for example). If I had directly pursued my original project of porting Arm to Unix-Xwindows, I would probably have had something workable by now, but all that would have happened is that Penni Sibun & perhaps Phil Agre would have played with it for 5 minutes & then put it aside, without every figuring out what the point was.

Perhaps I'm too steeped in the savage rhetorical traditions of linguistics, but it seems too me that PCT papers are much too `nice' (especially Rick Marken's - Bill is a bit nastier in his cataloguing of blunders) in their treatment of the opposition. Perhaps everyone should go and (re-)read Chomsky's review of Skinner's `Verbal Behavior' to see how it's done. For example, in `Control as Fact and Theory', Rick says that Fowler & Turvey argue against control theory, but doesn't give any indication of how thoroughly they misrepresent it, or how egregiously stupid their misreading is. In a linguistics paper their would be at least a page devoted to this theme (not, of course, saying `they're being stupid', but presenting evidence from which the conclusion follows immediately). But maybe the relatively rigid conventions that psych papers are written by make this harder to do.

Avery.Andrews@anu.edu.au

Date: Tue Feb 02, 1993 5:47 pm PST
Subject: demo of linguistic nastiness

[Avery Andrews 930203.1102]

For example, here's how a Chomskyan linguist might try to inject PCT into the `motor-action' controversy (without the references). Substantively, this needs a lot more work, but the rhetoric is typical linguistics (it might work on Penni's crowd as well, given enough substantive backing). The basic stragic themes are

a) minimal talk of revolution

b) purport to present common sense solutions to concrete problems

c) present evidence from which it is a short deduction that the opposition are shortsighted and seriously confused.

The study of motor control is currently split into two schools, `action theory' (Kugler, Turvey, ..) and `motor programing theory' (Schmidt, ...). Here we will argue that neither of these approaches shows much promise of providing concrete insight into what skills actually are or how they are acquired, and will propose a third approach, PCT, which combines the conceptual strengths of action theory and motor programming theory without their most obvious drawbacks.

The older of these approaches, motor programaming theory, began as the doctrine that many or even most actions were internally represented as sequences of motor commands to the muscles: when an action was to be initiated, the program was activated, and the required movements would follow. However, it was quickly realized that this kind of approach could not work: motor programs needed `parameters' that could specify variable aspects of the desired performance (e.g., size, for handwriting), as well as access to reafferent feedback to adjust the performance to variable aspects of the current situation. But there appears to be very little in the way of substantive proposals as to how this more more general notion of motor program might be implemented in the nervous system. Therefore, although the general claims made by motor programming theory are undoubtedly true in some sense, not much insight into the structure of ordinary behavior has followed from it: we still don't know how to specify a motor program for reaching for a cup of coffee, let alone lifting it, or getting the contents into rather than onto our faces [good one, Rick].

The alternative, action theory, is in many respects a major innovation in psychology. Motor programming theory uses the vocabulary, and perhaps depends on the concepts, of traditional `event-based' stimulus-response theory, whereby the sensory input is thought of as being divided into discrete `stimuli', and action into discrete `responses'. This is obviously a serious oversimplification, which action theory quite rightly rejects, emphasizing the fundamentally continous nature of both perception and action. Action theory also correctly recognizes the body and the environment as essentially involved in the organization of action: muscles on their own, for example, don't produce useful results. They only produce forces, which may or may not produce useful results, depending on the dynamics of the environment.

But, incomprehensibly to us (as well as to motor programming theorists), action theorists appear to be opposed in principle to saying anything substantive about the internal structures, essentially those of the nervous system, wherby complex activities are conducted and learned.

---fn
A possible rationale for this is fear that if people feel free to talk about the nervous system, they will feel free to try to explain the nature of behavior in terms of properties of the nervous system in isolation (Kugler et al 1980:10). We agree with the concern, but not with the strategy of ignoring the nervous system.

This makes it very unlikely that this approach as it stands will be able to offer much insight into the nature of complex skills and how they are acquired.

PCT on the other hand, provides an approach to the study of behavior which is inherently dynamic, interactive and `situated' (virtually all behavioral results are achieved through continuous, closed-loop interaction with the environment), but in which one can also make straightforward and commonsensical proposals about what kinds of internal structures (presumably neural circuitry, for the most part) make it possible for organisms to act, and to learn to act.

PCT is a closed-loop theory of behavior, so it is necessary to explain why it is different from other closed-loop theories, such as that of Adams (1968, 1972, 1977), that have been generally abandoned. We argue that these theories suffer from a combination of outright misunderstandings about how closed-loop systems actually work, with arbitrary and in fact incoherent limitations on what feedback is. With the limitations and misunderstandings removed, feedback theory will be seen to be a far more powerful tool for analysing behavior than has generally been realized. It also leads to fundamental revisions in the foundations of psychology - `stimulus' and `response', for example are no longer basic notions, although they may still be useful for certain kinds of situations.

.....

well that's all for now, folks.

Date: Tue Feb 02, 1993 7:34 pm PST
Subject: output blunder, CYA

[Avery.Andrews 930203.1422] (Bill Powers (930201.1900))

>>But is this really a blunder? After all, if an `output' is
>>used as a reference level for a control system, you do get a
>>regular result.
>
>How are you going to use an output for a reference level? I think
>you might have meant to say "if a reference level is set for an

What I meant was this. Suppose the system wants to move its hand up, then down. A higher-level system can then send a rising-falling waveform as a reference-level to a lower level height-control system.

This waveform can be thought of as a `command', which thanks to the lower level control systems, will indeed deliver a consistent result in spite of the variations in the environment.

>>I think it's important not to spend too much energy on
>>peripheral matters where PCT expectations might not be borne
>>out -- if you insist that the jab is done via kinesthesis, and
>>somebody proves that it isn't, then you've lost a lot of
>>credit.
>
>The whole point of making models is to make them testable, so
>they CAN be disproven if nature so chooses. I'm not out for the
>kind of credit you're talking about (speaking so cautiously and
>generally that you don't risk being wrong). I think that sort of

I'm not advocated CYA, but rather putting the maximal effort on the most rewarding points. If you've got `real psychologist's' attention for ten seconds, it's important to make the most of it, which means dwelling on the areas where PCT is strongest, and the opposition weakest. I don't think that boxing jabs and similar very high speed actions are one of those areas.

Avery.Andrews@anu.edu.au

Date: Tue Feb 02, 1993 9:13 pm PST
Subject: demo of linguistic nastiness, error control

[From Rick Marken (930202.2100)]

Avery Andrews (930203.1102)

>For example, here's how a Chomskyan linguist might try to inject PCT
>into the `motor-action' controversy (without the references).

Your example article is wonderful. If I were as good at this writing business as you are, your description of the relationship between motor program, action theories and PCT would have been the introduction to my paper on "Degrees of Freedom". Very nice work Avery.

Mark Olson (930202) --

> Now I am curious, would you not agree that its a game of semantics once
>you say that this error that is controlled is a perceptual signal for
>another system?

Probabaly. "Error" describes the functional significance of a particular signal in a control loop. This signal could be the object of control by another control system, in which case it would be OK to say that "error" is being controlled.

>Basically what am asking is whether you agree that this is semantic
>quibble since I believe we have the same model in our heads. I understand
>that error is controlled Because it IS represented as a perceptual
>signal....but it's still error nonetheless.

Sounds good to me; error can be controlled; but when it is, it must be represented as a perception (of the error signal) in a control system that has its own error signal that indicates the discrepency between what it intends and what is the actual level of perceived error.

Right? Best Rick

Date: Tue Feb 02, 1993 9:43 pm PST
Subject: Linguistic nastiness

[Avery.Andrews@anu.edu.au] (Rick Marken (930202.2100)) (Avery (930203.1102))

I'm glad you liked it! But of course the flashy rhetoric can at the most get people's attention for about ten seconds, & has to be followed immediately by some concrete results they can relate to.

My rhetorical component can probably spit out another section on how the `locality condition' in the `technical definition' of feedback prevents there from being a non-mystical theory of coordinative structures, at which point some actual work will have to be produced, along the lines of yours. Your papers make the point, I think, but not in a sufficient lurid way to make the kind of impression needed (the punters don't see something that's enough like ordinary behavior being produced). I appreciate them a lot more after my recent dip into the motor control literature, but people who haven't already looked at this stuff while cultivating an interest in PCT might well not get it.

One possibility would be tip-position control by a `tentacle', e.g. a light multi-jointed arm assumed to be underwater, so that gravity is irrelevant, and viscous effects dominate inertia, so that the dynamics is not such a problem. The `coordinative structure' being compensation at some joint angles for disturbing forces applied at others. (On the general subject, I've come across a number of references on arm positioning and coordinative structures by a guy called Cruse, mostly in Biological Cybernetics.)

Returning to the putative paper outline, producing (a) a sample coordinative structure (b) a system that acquires a strategy (such as learning that you have to turn away from a noxious stimulus before putting on a high-speed dash to get away from it) might be enough substance to go with.

But I would need to learn a lot more about coordinative structures, I think, though what I've read so far is confusing enough to make me suspect that there must be lots of people out there who would be interested in a common-sensical alternative.

And I do not joke when I advise people to read Chomsky's review of Verbal Behavior.

Avery.Andrews@anu.edu.au

Date: Wed Feb 03, 1993 6:23 am PST
Subject: Devilish misc.

From Greg Williams (930203) Bill Powers (930202.1200)

>Only by determining the actual muscle
>spring constants and force-velocity damping could we demonstrate
>that Bizzi's constants can't be those of the muscle, and must be
>set by parameters of a control system. Could this be done using
>data from the literature?

Somewhere (I haven't been able to find it) there is data purporting to show that trajectories altered by disturbances don't go the way they would if there was end-point control. Maybe Avery can find it among the Bizzi (I think) papers and evaluate the claims.

>So inverse kinematics still figures into the model. The central
>computer must pick a desired path, compute the sequence of joint
>angles necessary to produce that path, and emit the appropriate
>waveforms of signals to the joint-angle positioning systems. In
>the control-system model the inverse kinematic computation is not
>carried out; the same waveforms are generated using higher-level
>control systems, both kinesthetic and visual.

I'm not so sure. Much of what Bizzi and his colleagues write is vague and obscure (at least to me) -- one reason I suggested asking him directly for clarifications. If he believes what PCTers think he believes, OR if he actually is closer to PCT ideas than PCTers think, EITHER WAY, I can't see any harm in direct communication. The old racist ploy of deliberate isolation from "the devils" (in part, to make it easier to maintain the "devils" label, which is more difficult when there is personal interaction with the "devils"!) has no place in science, in my opinion.

>A month or so ago I sent Bizzi a letter containing the rejected
>Science article, the writeup of the Little Man program, and a

>disk with the program and source code on it. I have not even
>received an acknowledgement of receipt.

So why not phone him? The paper affords a perfect entre for direct communication. You could ask him what he thinks of your ideas and what is fundamentally divergent in his and your models, rather than guessing about (devilish?) motives.

>>In abandoning strict end-point control, beware the infinite
>>regress: controlling for "being in the process of" "being in
>>the process of" "being in the process of"... "being in the
>>process of" going to Paris! How do you decide where to draw the
>>line?

>That's just playing with words. "Going to Paris" consists of
>setting and accomplishing numerous subgoals, each of which
>requires control to be realized in a variable environment. Some
>of those might also be processes, but most will not be (like
>reaching for the telephone).

Now I think you are retaining the notion of end-point control (i.e., error between being where you are now (not Paris) and being in Paris) results in "setting and accomplishing numerous subgoals." Yesterday I had thought you were proposing reference signals for "being in the process of" thus-and-so, and I still think they are unnecessary.

>>I think
>>it is more parsimonious to hold to end-point control, with
>>transient conditions almost always present for the higher-level
>>(slower time-scale) processes.

That was the main point I wanted to make: that, in general, higher- level goals are NOT satisfied, but only in the process of BEING satisfied. Why do you want to maintain that goals are practically always satisfied over time? I suppose it is to avoid having reorganization come into play all of the time! But you could postulate that the time scale for reorganization beginning depends on the inherent time constants of the levels where goals aren't satisfied (longer until reorganization starts

>But then you've never bought into the idea of levels of control,
>so you're sort of stuck with end-point control, aren't you?.

You misunderstand. I haven't seen sufficient evidence to justify YOUR vision of the way peoples' control structures operate as either the ACTUAL way or the only POSSIBLE way. But you might be right. Your proposal is simply too audacious for me to buy, given the current evidence from neurophysiology and psychology, which is meager. However, what evidence there is (especially from studies of "simple" organisms) suggests to me that some kinds of hierarchies of control are likely. But "higher-level" circuits might sometimnes alter gains of (especially, inhibit or disinhibit) "lower-level" circuits instead of altering the latter's gains.

What would really get away from strict end-point control (no matter how sophisticated) would be adaptive/model-based control.

>Hope Cam arrives at a suitable end-point. That sounds like a very nastyaccident.

The doctor stomped on his arm until the bone alignment looked pretty good -- didn't cut it open to put in a plate. But he might later, if the bones don't knit. The biochemical control circuits are in charge now.

>Avery Andrews 930203.1015

>Asking the devils what they mean is not going to work, since (a) they
>won't answer

How do you know until you've tried? Very few attempts to communicate with "devils" have been made to date by PCTers.

>(b) it will just reinforce the image of PCT-ers as an
>annoying group of people who niggle about arcanities of formulation.

Well, aren't we? :>'|= (The apostrophe is a chip-on-shoulder.) Seriously, if that image is already there, only personal interactions are going to dispel it. (Believe me, it does help to interact with, for example, Rick, personally -- if everybody did, we probably wouldn't joke about "despite what they say on the net about him"!)

>Perhaps I'm too steeped in the savage rhetorical traditions of
>linguistics, but it seems too me that PCT papers are much too `nice'
>(especially Rick Marken's - Bill is a bit nastier in his cataloguing
>of blunders) in their treatment of the opposition.

Gee, and then maybe we can get a reputation like those "savage" linguists (and anthropologists, too, I might add) have among those in "harder" scientific fields. I think many would join me in thinking that rhetoric is a poor solution for frustration. One more time: (1) Do the devils have error signals? If so, go to (2); otherwise, give up. (2) Can PCT help them reduce the error signals? If so, go to (3); otherwise, give up. (3) Publish the PCT aids to the devils' error signals. (4) Have patience, in any case. Is there really ever any point to advocating rapid revolution other than personal aggrandizement?

As ever, Greg

Date: Wed Feb 03, 1993 6:40 am PST
Subject: effective argumentation

[From: Bruce Nevin (Wed 93023 08:49:30)]

(Avery Andrews 930203.1102 & Wed, 3 Feb 1993 16:38:51 EST) --

Excellent! I like this a lot.

>And I do not joke when I advise people to read Chomsky's review of
>Verbal Behavior.

Seconded. It was available in a Bobbs-Merrill reprint for many years, probably still is. Chomsky is a master of debate and a devastatating polemicist. Bob Ingria, a former student of his and Morris Halle's, who has been here at BBN for a number of years, tells me "if you want to talk with Chomsky, wear boxing gloves." He has attracted like minds (or only like minds have survived in his company). Any would-be revolutionaries should study his successful techniques.

Recommended:

Botha, Rudolf P. 1989. Challenging Chomsky : the generative garden game. Oxford, UK ; New York, NY : B. Blackwell.

_____. 1980. The conduct of linguistic inquiry : a systematic introduction to the methodology of generative grammar. The Hague ; New York : Mouton Publishers.

_____. 1973 The justification of linguistic hypotheses; a study of nondemonstrative inference in transformational grammar. With the collaboration of Walter K. Winckler. The Hague, Mouton.

Also, barbed tongue in lacerated cheek, Paul Postal (an angry, now no longer young, man), "Advances in Linguistic Rhetoric", a few years back in the Topic/Comment section of (help me out here, Avery, I can't find my reprint or the reference).

On another topic, BBN just announced that it will lay off 300 people over the next 6 months, most of them from my division, so it more than ever behooves me to be productive in meeting reference perceptions held in common with others who are working to get certain new products out the door, seeing that they meet customers' needs and expectations, etc. This has delayed my response re language. Since I had said all that I want to say previously, but evidently had not communicated it, formulating appropriate response takes time and care that I have not been able to bring to it. I have no major quarrel with what has been said here, as far as it goes, only it does not go far enough: it omits the essentially social character of language. The human control system does not create language anew, but rather encounters it pre-existing as a social reality, and re-creates it in that encountered image. That is how it is that language can be described as a natural object (indeed, a mathematical object) whose evolving structure in part reflects perceptual control as its byproduct, to be sure, but also in strong measure constitutes an aspect of "boss reality" guiding (limiting, disturbing) perceptual control.

Here's one aspect of the problem: the extent to which we live in a world of the imagination, strongly colored by our verbalizations, is I think insufficiently appreciated. The grease of social agreements lets us get away with ignoring (or approximating, by convention) a little raspiness in our feedback through the physical environment.

Bruce bn@bbn.com

Date: Wed Feb 03, 1993 9:13 am PST
From: Marken

MBX: Marken@courier4.aero.org
TO: * Dag Forssell / MCI ID: 474-2580
Subject: Meeting

Dag, I don't know if I will be able to make the meeting tommorrow. Linda has another meeting she must go to and Lise may have transportation demands that I must cover (she's out of school thanks to the year round school calendar that seem to let the kids go to school even less than they did before).

Anyway, if I can make it I'll be there rooting for you. If not, say hi to Todo for me if he is there.

Good luck. I know you'll do great.

Best Rick

Date: Wed Feb 03, 1993 11:10 am PST
Subject: Re: Effector-result relationship

[From Oded Maler 030293] [Rick Marken (930202.1200)]
*
* Martin Taylor (930202 13:45) --
*
* >But "flinging it into the air" is a common result of picking up something
* >you thought to be heavy, but isn't.
*
* I have NEVER had this experience with a suitcase or any other
* lighter than expected object. But thanks for the tip; I'll
* watch out if I'm ever checking in a very large, very light
* suitcase in Canada.
*
* >Likewise, failure of the initial
* >lift to move something off the ground is quite common when something is
* >heavier than expected.
*
* This only seems to happen when the suitcase is so heavy that
* some extra purchase (like a dollie) is needed to lift the
* thing at all.

Not necessarily, you may "estimate" that something is light enough so that you can pick it up when you are in some position, and than discover that you should approach it from onther position, and maybe use another combination of muscles in order to achieve the higher-level perception of lifting it up. The feed-back is taking place in the higher level system that sets reference signals for "pick up with finger muscles" and when this fails it sets references for "pick up with shoulder muscles". Servoing at this level is much more obscure than servoing at the quantitative level of "If you can't beat it - smash it".

[Bill Powers, on controlling long-term goals, etc.]

At the higher-levels there might be some advantage in shutting down feed-back pathes. It is maybe a question of span of attention. If you start contemplating about the meaning of life, death and all that each time you make a coffee, the outcome is not very efficient (from personal experience). In the army (and not only there) , soldiers and commander are usually praised for "sticking to the goal" (I'm not sure I'm translating correctly from Hebrew). Which means controlling very well at some lower-level, but maintining soe higher-level reference fixed (which means not to have feed-back at one level above). On the contrary, we say that people are "opportunists" or too dispersed, unconcentrated if they allow feed-back at certain levels (e.g., above ideology choosing). Trying to clarify to myself how these and other phenomena can make sense in PCT terms, I am more than before convinced that the location of (the perceptions of) "the same objective external variables" within the hierarchies of two different individuals might be *very* different. One man's system concept is another man's sensation (and of course, only words, those poor approximate illusions, give any reason to identify these completely different perceptions, which are controlled for on completely different time-scales, with each others). I have the feeling that I did not make myself clear on that (but the text contains self-justification of it).

Anywau, I just submitted today a paper I was working on for a month, and I have the famous post-submission effects..

--Oded

Date: Wed Feb 03, 1993 12:59 pm PST
Subject: Bizzi et al; qualitative movements; end-point control

[From Bill Powers (930203.1000)]

Starting on Sunday, Feb. 7, I'll be gone for 10 days: a trip to Green Valley AZ for by father's 93rd birthday on the 8th, then to San Diego to see relatives, then wandering around the southwest until we feel like coming home. Back approx. the 17th, so everyone who has projects going with me can do something more useful until then.

Avery Andrews (930203.1015) --

That is a wonderful start on the Big Article. I vote that you be principal author, with others of us contributing bits and pieces to flesh it out. Are you willing? I know you will have other committments, but even if it goes slowly it will be worth taking the time. How about a mid-summer target date for completion?

>My understanding of Bizzi is that he's aware of the possibility
>that feedback is involved, but thinks that he's shown that its
>contribution for head-orientation movements is rather modest -
>about 10-30%.

We need to know how he figured that out. Has he deafferented the control systems so he can measure the real muscle parameters without the feedback being involved? Can anyone help on this?

>Perhaps I'm too steeped in the savage rhetorical traditions of
>linguistics, but it seems too me that PCT papers are much too
>`nice' (especially Rick Marken's - Bill is a bit nastier in his
>cataloguing of blunders) in their treatment of the opposition.

I was warned by a number of people that calling them "blunders" was too nasty. I think you can only be agressively nasty when you have a prominent position in the community and a big enough ego to go with it, as well as a tendency to forget that your opponents are probably perfectly nice guys face to face.

(930203.1442) --

>Suppose the system wants to move its hand up, then down. A
>higher-level system can then send a rising-falling waveform as
>a reference-level to a lower level height-control system.

Yes, this is probably how intentional movements are created (like directing an orchestra). But again the qualitative-quantitative problem arises. A higher control system must always emit an output signal of a specific magnitude, which says HOW FAR up and down the hand will move. Even when you don't consciously intend to produce a quantitative effect, the actions of control systems are always quantitative. There's no way to tell an output function "just move the hand up a little bit." That means nothing to a neural network; the command has to be turned into a specific movement, and the signal always has a specific quantitative magnitude. The only way for the nervous system to be vague is at the symbolic levels.

>This waveform can be thought of as a `command', which thanks to
>the lower level control systems, will indeed deliver a
>consistent result in spite of the variations in the
>environment.

Not if the command entails a relationship to something else that can change. You can't issue a command to "move your hand to the vicinity of the cup" if someone is moving the cup. And if you're able to grasp the cup anyway, this means that there's a control system involved even when there isn't any disturbance. Moving "up" means moving in relationship to something else. If the frame of reference moves, the command won't be obeyed properly -- unless the relationship to the frame of reference is being perceived and controlled.

I think the principle here is that if disturbances CAN happen, the nervous system must be designed to assume that they WILL happen. So even issuing a command to move up and down means that a reference-perception for an up-and-down event is in place, and any disturbance will be resisted.

Greg Williams (930203) --

>Now I think you are retaining the notion of end-point control
>(i.e., error between being where you are now (not Paris) and
>being in Paris) results in "setting and accomplishing numerous subgoals."

End point control involves comparing what you're experiencing with what you want to experience and turning the error into an action that will make the error smaller. When the goal is positional, an error signal can be translated more or less directly into changes of reference signals that direct motor efforts, in a way that will in fact start reducing the error. This is how I have thought of configuration control, which is truly end-point control. There is no control of the path between end-points; the error always drives the system as fast as it will go toward the state of zero error, the path being determined only by system and environment dynamics.

This sort of control is practical only when a configuration error can create a vector effort aimed at the required goal state. The circumstances have to be very favorable for configuration control to work. Organisms with no levels of control higher than this can maintain bodily configurations, but can't handle intervening obstacles that might prevent moving toward a selected configuration. They just bump into the obstacles. They aren't smart enough to go around them. They can't even stop trying when errors are prevented from being corrected.

Rate of change control, or transition control, is a little smarter. With this level of control, which exists in the crowd program, the organism can turn away from the configuration-goal in order to keep its velocity going, and this will lead it around most obstacles. Velocity control has no end point.

The crowd program is, at its highest level, a spatial- relationship controlling program. It therefore can't control for a strategy, like pushing an obstacle out of the way in order to pass through an opening. Neither can it reason: if I take this nearest opening I will end up in a trap, so I'd better pass it up and look for another way. The pure relationship-controller simply heads in the direction that reduces the immediate relationship errors.

Going to Paris certainly entails controlling a configuration: being in a scene that one recognizes as Paris. But simply setting that scene as a reference signal and trying to correct the error by controlling efforts won't work. A true end-point control system would simply experience error -- it's unlikely that such a scene would ever become a reference signal because no end-point control system could correct the error.

One of the main reasons for my adopting a hierarchy of control is that errors between abstract perceptions and reference signals can't simply be routed to the muscles. That is too big a jump; the information in the abstract error signal isn't the right kind to be used in generating an appropriate muscle tension. If your checkbook balance doesn't agree with the bank statement, which muscle should you tense first?

It seems essential that high-level errors be converted into changes in lower-level reference perceptions in a series of stages. At each stage, the error is converted into a somewhat less abstract goal-setting. While I was trying to define a specific hierarchy, the main question was always "what is the next less general kind of goal that has to be changed in order to start correcting this kind of error?" If the error is in my bank balance, one thing I have to do is look for an arithmetic error, which involves making quantitative comparisons and correcting any that don't check out. If I could successfully do that much, the result would be to reconcile the bank statement with my checkbook.

Now, how do I correct a quantitative error? If I could change the number written in my checkbook to the correct number, that would accomplish the goal, so the error would be corrected, so the checkbook would balance.

How do I change the number I see in the checkbook into a different reference-number? One subtask is to remove the wrong number. I might do this by picking up an eraser and obliterating the number. So picking up the eraser and rubbing it over the number is a goal which, if successfully achieved, would remove the wrong number. Now how do I do one of those tasks, say picking up the eraser? If I could move my hand to it, grasp it, and lift, I would have picked up the eraser. How to accomplish that? I have to make effort-vectors appear. And to make an effort-vector appear, I have to feel muscle tensions.

As one goes down levels in such a hierarchy, the subtasks that are selected are not specific to the higher-level error. They're used to correct all kinds of higher-level errors. So there can't be any one simple path down the levels.

Now maybe that's not the only way to see the various levels of control, but it does satisfy the main criterion: that the adjustments of each level's reference signals can be based on the error signals of the next higher level, without any conceptual gaps. At each step, a subgoal is chosen which, if it were achieved, would be sufficient to move the higher-level perception closer to its reference signal. The step from one type of perception to a lower-level type is small enough that translating error into a change in subgoals can be done directly and appropriately.

>Yesterday I had thought you were proposing reference signals for "being
>in the process of" thus-and-so, and I still think they are unnecessary.

I've been trying to explain here why I think they're necessary. But beyond the logical problem of converting an abstract error into more detailed goal-adjustments, there's another fact that's just as important as the abstract principles involved. It's that people actually do control for the form of a process. This is easiest to see in simple situations, like kneading batter for bread. Doing this involves creating a long pattern of behaviors of the hands and the bread dough -- rolling, folding, squeezing, shaping, rolling some more, and so on (I don't bake bread, so don't laugh). This process simply continues, and if you're an expert it has a certain form that must be maintained as long as it's going on, until a higher-level control system decides that enough kneading has been done and it's time to go on to the next step.

While the process is going on, you don't say "Now I'm picking up the dough. Now I'm turning it over. Now I'm throwing it down. Now I'm folding it." That's how you would talk if you were perceiving only at the event level (assuming you could still talk). What you say is "I'm kneading the dough." If the immediate process-goal is to be kneading the dough, it is satisfied the moment you begin, and you simply correct errors that disturb the process, like the dough sticking to your hands, or momentarily forgetting where you are, or dropping the dough when you didn't mean to, and so on. The process control system corrects the errors by altering the subgoals of lower level as required to maintain the form of the process.

"Going to Paris" is an abstract goal. It doesn't mean just experiencing a certain visual configuration. We call such processes "a trip," and it is a pattern made up not only of the place, but of all the rituals and procedures involved in taking the trip. When you think of going to your office, you don't just think of sitting there in the office (although that's part of it). You know immediately how you're going to get there; the whole path springs to mind, with whatever details are appropriate to the weather or the season. The sense of this entire picture is the reference-signal.

You can't make your perception of the trip match the reference perception just by striking off toward the end-point. If you do that, you'll bang your nose on the door if you haven't fallen over a chair first. Going to the office entails a lot of little subgoals that have to be sequenced properly -- open the door and THEN walk through it. As long as all these details are occurring in the familiar pattern, you know that there is no error in carrying out the process of going to the office. If you were asked what you are doing, you'd say "I'm going over to the office." That's your description of the process that you've started and are maintaining RIGHT NOW.

>... in general, higher- level goals are NOT satisfied, but only
>in the process of BEING satisfied. Why do you want to maintain
>that goals are practically always satisfied over time?

I didn't say they ARE always satisfied over time. I just said that a wise person will select higher-order goals so they can be achieved immediately: don't aim to be a millionaire; aim to do all the things that accumulate money rapidly. That is much more likely to succeed, and you can satisfy that goal immediately. If you always set distant goals, you will have very little control. If you adopt a small enough time-unit, all errors in all control systems are in process of being corrected. But the process itself is not the goal, in that case. This is an example of using an inappropriately high level of perception to look at a low level control system, like characterizing all control actions as sequences, or programs.

The opposite error is to use too low a level of perception in characterizing a higher-level control process. One could object that nobody actually does something called "Going to Paris." All that people actually do is to make certain muscles tense while relaxing others. This was B. F. Skinner's kind of argument against purposive descriptions of behavior. He said that nobody has the purpose of mailing a letter. There are simply certain responses that cause the hand to grasp the letter, and then that cause the body to be propelled down the street, and then, if the environment is so configured, make the hand rise and put the letter in the mailbox. There is no actor in this event; it is simply that conditioning and discriminative stimuli happen to create a chain of events that, this time, ends up with the letter in the mailbox. If the environment had been different, the letter (presumably) could have ended up in a trash barrel or a fire- alarm box, or sticking out of your ear.

If you adopt the time-scale appropriate to the speed with which any level of control can correct its errors, the process of error correction occupies the shortest distinguishable time unit and can be ignored. The kind of process control I am talking about is concerned with variables that change very slowly in comparison with the time it takes to correct an error in the process. I'm talking about controlling perceptions, where the perception itself represents the presence of an ongoing process.

>You misunderstand. I haven't seen sufficient evidence to
>justify YOUR vision of the way peoples' control structures
>operate as either the ACTUAL way or the only POSSIBLE way.

In the binary choice between sufficient evidence and insufficient evidence, I suppose you're right. I think I have some reasonable ideas, and many informal examples that seem to fit them. But that falls between "right" and "not right" so I suppose it's in territory that you don't acknowledge as meaningful.

>But you might be right. Your proposal is simply too audacious
>for me to buy, given the current evidence from neurophysiology

>and psychology, which is meager.

There is quite a bit of evidence from neurology that we perceive sensations, configurations, transitions, relationships, sequences, and probably some other classes of experience that are in my proposals. But all this evidence rests on our ability to perceive such things directly -- all that neurophysiology can tell us is what part of the brain has been damaged. The behavioral or perceptual correlates of that physical damage can be understood only subjectively -- we see that another person seems unable to perceive or control something that we ourselves perceive and control easily, so perhaps that has something to do with damage to that part of the brain's geography. Neurology will never, by itself, tell us what any part of the brain is for.

>But "higher-level" circuits might sometimnes alter gains
>of (especially, inhibit or disinhibit) "lower-level" circuits
>instead of altering the latter's gains.

No objection to that, but this leaves open the question of what variable the higher system would be controlling by altering the gain in a lower system. In order to model this, we would have to propose some perceivable aspect of the lower systems' behavior that the higher system compares with a reference signal, and controls by varying the gain of the lower system. One idea might be oscillation. If the higher system detects oscillation, it lowers the gain of the lower system until the oscillation disappears. I've mentioned this mode of control many times, but have never actually modeled it.

It's not necessary to have a separate function for inhibiting and disinhibiting lower control systems (I've spoken about this before, too). All neural control systems are necessarily one-way, because neural signals can't go negative (and muscles can't push). A two-way control system must actually contain two comparators, one handling signals that are positive for one direction of output, and the other handling signals that are positive for the opposite direction of output. In all instances of two-way control, the actual neural control systems must occur in balanced pairs, both handling positive signals but perceiving signals having opposite meanings and producing outputs having opposite effects.

As a result, a two-way neural control system is turned off if both sides of the balanced pair receive zero reference signals, given the arrangement in the canonical diagram where the feedback signals are inhibitory. If both reference signals are zero, then no matter how large the perceptual signals are, inhibiting a neuron that receives zero excitatory input will produce no error signal. The dual system is effectively turned off.

>What would really get away from strict end-point control (no
>matter how sophisticated) would be adaptive/model-based control.

I wish you wouldn't identify "control" with "end-point control" End-point control is only one kind. Controlling the rate at which you crank a windlass is not end-point control. Controlling a continuing action like dribbling a basketball or fanning your face with a newspaper is not end-point control. Maintaining a specific distance from a randomly-moving object is not end-point control. Speaking to someone in a manner you (at least) perceive to be civil is not end-point control. Trying to be a good scientist is not end-point control. You seem to be forgetting that control is not just error-correction -- it maintaining some perception in a particular state, and that perception can involve dynamic patterns as well as end-points.

While adaptive control and reorganization and gain control and other subjects are important, I think it is premature to worry about them. I'm still trying to get people to see the significance of ordinary control, the simplest kind. Most people, even those interested in PCT, seem to be satisfied with noticing a few phenomena in other people's behavior that verify the principle, but have very little interest in reinterpreting their own everyday experiences as control of perception instead of just "doing things." People tend to back off from the task of reinterpreting themselves to themselves. Even in our group, everyone has some idea they're not willing to look at very closely -- the ideas that are "obviously right." So it goes.

Bruce Nevin (930203.0849) --

>Here's one aspect of the problem: the extent to which we live
>in a world of the imagination, strongly colored by our
>verbalizations, is I think insufficiently appreciated. The
>grease of social agreements lets us get away with ignoring (or
>approximating, by convention) a little raspiness in our
>feedback through the physical environment.

How do you, personally, know there are any social agreements?

Best to all,   Bill P.

Date: Wed Feb 03, 1993 2:18 pm PST
Subject: Big Article, Reticent Reinterpretation

[From Rick Marken (930203.1300)]   Bill Powers (930203.1000) --

>Avery Andrews (930203.1015) --

>That is a wonderful start on the Big Article. I vote that you be principal
>author, with others of us contributing bits and pieces to flesh it out.

I second the motion. Avery knows more about motor control after eight weeks than I know about it after eight years (I hope he doesn't expect me to understand linguistics as thoroughly as he now understands motor control).

I think you (Avery) have an excellent grasp of the issues and you are looking at them from the right perspective;that of the interested bystander who can say (without cant) "gee, if you do it that way isn't there the possibility that it will break if anyone breathes too hard?"

I would be happy to try to supply the supporting research (if it can be done with minimal hardware) or modelling that you think might be useful for the article.

I hope you can do it.

>I'm still trying to get people to see the
>significance of ordinary control, the simplest kind. Most people,

>even those interested in PCT, seem to be satisfied with noticing
>a few phenomena in other people's behavior that verify the
>principle, but have very little interest in reinterpreting their
>own everyday experiences as control of perception instead of just
>"doing things." People tend to back off from the task of
> reinterpreting themselves to themselves.

That's the fact, Jack!

Best Rick

Date: Wed Feb 03, 1993 3:11 pm PST
Subject: dances with devils

[Avery.Andrews 930204.0940] Greg Williams (930203)

>So why not phone him? The paper affords a perfect entre for direct
>communication. You could ask him what he thinks of your ideas and what

Because if his articles are obscure, what he says in a phone-call from out of the blue will be even more so, & the entire experience is likely to be extremely embarrassing. If the writing looks obscure, you just publish something claiming not to understand what it means, & leave it to the author to clarify, if they are able and willing.

>>Asking the devils what they mean is not going to work, since (a) they
>>won't answer
>
>How do you know until you've tried? Very few attempts to communicate
>with "devils" have been made to date by PCTers.

Because I am one myself and I therefore know what they're like. I think a number of conclusions can be drawn from (a) typical referee responses to PCT papers (b) the way the communications with Randy Beer and David Chapman proceeded. I see no reason to think that unsolicited phone calls would turn out any better.

>Seriously, if that image is already there, only personal interactions
>are going to dispel it.

Maybe, but not over the phone.

>Gee, and then maybe we can get a reputation like those "savage"
>linguists (and anthropologists, too, I might add) have among those in
>"harder" scientific fields. I think many would join me in thinking
>that rhetoric is a poor solution for frustration.

Well, yes, rhetoric (e.g., presentational techniques) has to be used carefully. I think it's good to use it/them to facilitate understanding of what you're saying, but wrong (and probably just plain ineffective, in this case) to use them to create belief in what your saying. What the rhetoric amounts to is (a) using a bit of vivid language to heighten (re-)perception of absurdities that people have become habituated to (b) selecting what you say to address the immediate concerns and predelictions of your audience.

For example, Dag Forssell's approach of starting out with a discussion of what a good theory is would fail utterly with the kinds of people I hang out with, because we think we eat six theories for breakfast every morning (and everybody's ego is based on the idea that they're just about the smartest person they know). We may in fact be completely deluded in imagining that we know a good theory when we see one, but it's useless to tell us this explicitly. Ultimately, my crowd are the same as his: they will like it if it helps them solve problems they are already interested in, otherwise not. But the initial approach has to be different.

Avery.Andrews@anu.edu.au
Date: Wed Feb 03, 1993 3:14 pm PST
Subject: Big Article

[Avery Andrews 930204.0919] Rick Marken (930203.1300)

I suspect that all those years of trying to figure out Chomsky give me a bit of an edge in sifting through conceptual issues, but there's still *lots* I don't know about motor control. I'm happy to proceed more or less as I am, & we'll see what comes out.

Bill Powers (930203.1000)

>We need to know how he figured that out. Has he deafferented the
>control systems so he can measure the real muscle parameters
>without the feedback being involved? Can anyone help on this?

What he did is in the Bizzi et. al. (1978) article that I've mentioned a few times. I can't tell how valid it is, since it depends on engineering rather than the sort of conceptual analysis that I seem to be good at. At the moment my inclination would be to take a neutral attitude towards Bizzi. I think the relative role of mechanical vs. `reflex' elements in muscle stiffness is a fairly technical issue, best left to people like Rack., Houk & Rymer. `Coordinative Structures' and `Distributive Compensatory Responses' are what look like the rewarding targets to me.

In print, I think it would be better to call the blunders `misapprehensions'.

Avery.Andrews@anu.edu.au

Date: Wed Feb 03, 1993 5:52 pm PST
Subject: devils, angels

[Avery.Andrews 930204.1224]

Today's find in the library is:

Whiting, H,T.A (ed) _Human Motor Actions: Bernstein Reassessed_, North Holland.

Bernstein is one of the culture heroes of Bizzi et. al. on the one hand, and Kugler, Turvey, et. al., on the other, and appears to have been a major Right Thinker, having, for example, a very nice feedback diagram (pg. 358, also pg.

130 of N. Bernstein 1976 _The Coordination and Regulation of Movements_, Pergamon press).

The devil's bib entry is Kugler and Turvey `An Ecological Approach to Perception and Action' pg. 373-412, esp. pg. 391-392, where they criticize Bernstein's and everybody else's conception of feedback control on various grounds, including what appears to be the idea that it involves `an orderly sequence of symbol strings (the representational format for the quantities and the commands)'. They appear to be making the event-based blunder and several others besides, and to be criticing feedback control in general on the basis of somebody's proposal for a symbolic `motor program' to control walking (due in its original form to MacKay, W.A. (1980) `The motor program: back to the computer', Trends in Neuroscience 3:97-100). I'm getting the impression that part of their strategy is to criticize actual proposals on the basis of the most bungled derivatives of those proposals that they can find (a tactic we should be careful to avoid).

On the other hand, in the next section, they make what strikes me as a sensible argument that processes that may look like feedback control w.r.t. a set-point aren't necessarily so, and that respiration rate-stabilization in fact isn't.

Then there's an interesting-looking argument by G. Hinton arguing that it is in fact a good idea to precompute torques, and discussing various features of muscles that make this easier to do than you would think (for example, the viscosity properties of the force-velocity relations give you what is in effect instant feedback control over the velocity of a limb: if the velocity is less than expected, the force exerted by the muscle will be higher, so the velocity lag will be corrected. If Bill Powers and Greg Williams haven't thought about this article, they ought to.

Date: Wed Feb 03, 1993 7:48 pm PST
Subject: picture proposal

Here's a thought which people may not like, but here goes anyway. The idea is to redo the ECS diagram to emphasize the `bushiness' of perception & action, and also to allow more explicitly for compensation. The redone diagram is supposed to go like this (apologies if someone has already done something like this & I've forgotten about it):

```
                          |                  |   |
                          v                  |   |
                     ------------            |   |
       ------------>| Comparator |--------   |   |
      |              ------------       |  |   |
      |                                 v  v   v
  -----------                          -----------
 | Perceptor |                        | Modulator |
  -----------                          -----------
   ^  ^  ^                             |  |  |  |
   |  |  |                             v  v  v  v
```

The perceptor is in general a leaky integrator, perhaps of various kind of energy from the environment, perhaps of neural currents. The modulator does the `feed-forward' work, such as, (a) converting a step-change in high level reference level into a gradual change in a lower level one (b) altering the nature of the output to compensate for an anticipated disturbance (e.g. tensing agonist and antagonist muscles preparatory to trying to catch something heavy) (c) choosing an appropriate strategy to push the perception back to the reference level (e.g., to get the finger tip closer to the target, it might be appropriate to flex or extend a given joint, depending on circumstances).

All of this stuff is already in the model, I think - the idea of putting it into the picture in this way is to make this fact harder to miss.

Avery.Andrews@anu.edu.au

Date: Thu Feb 04, 1993 4:29 am PST
Subject: control article

[Hans Blom, 930204]

What is the fastest way to get your spaceship to Mars? The solution is well known, although impractical: Apply full thrust until you are at the exact midpoint of the trip, turn your ship around and apply full thrust again, braking until you arrive at Mars with zero speed. This is an example of what is called 'bang-bang' or 'minimum-time' control, a control paradigm quite different from the 'stabilizing control' that is mostly discussed on CSG-L. Features of bang-bang control are: 1) outputs are either zero or at their limit, 2) the only important parameters are the times at which outputs go from zero to maximum or from maximum to zero, 3) in general, it is quite difficult to find optimal values for those times, and 4) for long periods of time (between the decision points) it may seem to an outside observer that control is absent, because nothing changes, because there is no modulation of the outputs.

This is, partly, the contents of William S. Levine and Gerald E. Loeb's paper 'The Neural Control of Limb Movement' in the December 1992 issue of IEEE Control Systems. Does the organism use bang-bang control? No. 'The experimental data ... show a substantial deviation from the optimal control model'. Why is that? Partly in order to protect the organism: 'the feedback from the joint sensors, while certainly present, would be too late to prevent injury if a human jumper tried to perform a mathematically optimal [i.e. top performance] jump'. Partly because 'it is important for both biologists and control engineers to remember that the control systems that have been invented to date are almost certainly a meager subset of all possible types of control and even of all control methods used in biological systems. Thus, the study of biological systems should not be confined to testing whether their performance is compatible with control schemes invented to date but must include detailed examination of their inner workings to discover new types of control'.

Some type of stabilizing control is needed in all cases where full-time control relative to a setpoint cannot be relinquished even for a moment. But stabilizing control is incompatible with top performance, such as in sports. In maximum height jumping, only the maximum height of the jump is important, not the full trajectory. In the Mars rocket, the output resources are used at 100% capacity during 100% of the time; the only control decision is to find the exact point in space-time of the turnaround. Mathematically, due to the non-linearity of the problem, finding this point is generally intractible and therefore usually a matter of trial and error (search) or creative insight. In humans, finding the optimal decision points requires a considerable period of tuning and fine-tuning (training).

The authors pose more questions than they provide solutions: '... much more work needs to be done before the above suggestions can be called a theory'. Yet, in my opinion, this paper provides some insights into why stabilizing control, which works so well in ordinary circumstances, breaks down when maximum performance is required.

Date: Thu Feb 04, 1993 7:05 am PST
Subject: for Penni

[From: Bruce Nevin (Thu 930204 09:42:00)]

Penni, if you're still here, this looks like a natural for you:

Linguist List: Vol-4-66. Wed 03 Feb 1993. Lines: 125
-------------------------Messages--
Date: Wed, 3 Feb 93 14:21:31 EST
From: rambow@unagi.cis.upenn.edu (Owen Rambow)
Subject: Workshop: Intentionality and Structure in Discourse Relations

CALL FOR PAPERS

INTENTIONALITY AND STRUCTURE IN DISCOURSE RELATIONS

21 June 1993
Ohio State University
Columbus, Ohio, USA

A workshop sponsored by the Special Interest Group on
Natural Language Generation (SIGGEN) of the
Association for Computational Linguistics

TOPICS OF INTEREST: Over the last few years, discourse structure relations (often called "rhetorical relations") have been extensively discussed in the text planning community. Two of the best known text planning architectures, McKeown's TEXT and the ISI text planners, have explicitly and successfully incorporated the idea of a bounded set of semantically meaningful, domain-independent relations between discourse units. At the same time, computational work on text structure development and analysis has highlighted the need for intentionality (often called "communicative goals") as well. The precise relationship between the rhetorical and intentional types of knowledge is unclear. Making the issue even more difficult, the theoretical status and essential nature of rhetorical relations has never been clearly articulated, and while communicative goals have been linked with Speech Acts and intentionality in general, the precise territory of such goals has also never been defined. The goal of this workshop is to bring together researchers from different fields, including discourse understanding, discourse generation, and linguistic discourse analysis, and to debate and explore the issues involved. In particular, the workshop will address the following questions:

1. What is the evidence for the existence of rhetorical relations?

What types of rhetorical relations are there?
2. What is the evidence for the existence of intentions? What types of intentions are useful to identify for communication?
3. What is the precise relationship between these two types of knowledge? Do intentional and rhetorical relations perform different functions (though they may be related), or are rhetorical relations the realizations of intentions, or should rhetorical relations be discarded as simply a misconstrual of intentions proper?
4. How do rhetorical relations interact with representations of Speaker's and Hearer's beliefs and desires?
5. How are rhetorical relations used in discourse understanding? How are linguistic clues and world knowledge brought to bear?

Note that this is not a workshop on a particular theory of rhetoric, but on the theoretical foundations and implications of theories of discourse structure and intentionality.

FORMAT OF SUBMISSION: Submissions are sought that address one or more of the questions outlined above; they should be presented as position papers, with reference to the author's own work. Submissions should be by email (ASCII files) and should not exceed 2 ASCII pages (exclusive of references). Submissions should be sent to rambow@unagi.cis.upenn.edu. Authors without access to electronic mail should send submissions to:

Owen Rambow

Date: Thu Feb 04, 1993 8:49 am PST
Subject: Assorted miscalenia

[From Oded Maler 930204]

This is a collection of thoughts follwing a desparate attempt to catch up, clean one's mail box etc. I can't give exact temporal references.

{Avery Andrews, on the priority of work on motor control}

Amen. I think this is a concensus. You can speculate and tell anecdotes on higher levels, but the only chance to build models is from the bottom-up. I would go further and say that one should abandon motor-behavior of humans and concentrate on lower inverterbrates - as I suggested in an unpublishe manifesto "why should we build artificial worms" (soon to be rejected from IJCAI-93, and to become a serious contender for rejection record).

Concerning language: I think I'm a very verbal person and a heavy user of language and I'm very fluent in my native language at least. I used to be a journal editor back in school and than in the engineering school (that's why a lot of engineering math left no traces in my mind). Nevertheless I have not learned *about* language as a scientific discipline (unless you consider the stanfard CS curriculum on formal languages) so what I'm going to say is very professionaly unqualified (like, say, Ali speaking about dynamics). My feeling is that *most* of the formal linguistic research is a useless contemplation about hypothetical structures. It may have intrinsic beauty or complexity but its relevance to human use of language is like the relevance of formal logic to the way humans think. The latter (logic as a simulation formalism for human reasoning) was my entry point into research, and it is this diappointment from this approach that pushed me toward the bottom-up - and I speculate that this is also part of Avery's motivation. The only linguist paper I've read was Gross' "on the failure of generative grammer" (he told me some more on Chomskian nastiness) so I really make too-sweeping generalizations based on ignorance, but this reminds me of a beautiful quote from Kurt Vonnegut's (sp.?) Cat's Cradle which I quote from memory (will appreciate the exact quote) : "Beware people who have devoted all their lives to learn something and finally found themselves no wiser than before; they are full of resentment to all other people who have come to their ignorance the easy way."

About converting little-man to X-windows: if done, I promise to play with it more than the average nouvelle west-coast post-structuralist neo-phenomenologist would.

Kugler et al.: my impression from Kugler's presentation at Aix was mixed. His lecture gave me a very bizarre impression. Later I looked at his book with Turvey and I had an opposite impression. Also some of his comments in the evening discussion made sense. Unfortunately he left before I had the opportunity to decide finally to which basin of attraction my opinion is going to flow so I am still at the bifurcation point. But still when a non-mathematician judges such work he should bear in mind that "minimizing the Hamiltonian", "limit oscillator" etc. might really mean something very elementary to some people in the same sense that ECS, CEV etc. can be elementary concepts for you. This does not exclude the possibility that the math is ceremonial but it is not necessarily so.

CSG-meeting: travels not financed by tax-payers violate my current principles.

Programs: I think most of the people who talk about "programs" are completely unaware of some of the most dominant trends in theoretical computer science, namely the theory of concurrency. Unlike traditional computability (Input-Computation-Output), these theories deal exactly with processes that run in parallel and interact with each other. There is an enormous literature on it and about various problems associated with it. For example, in sequential programs, when you take one sub-procedure, and replace it by another which computes the same output but more quickly, the overall meaning (behavior, function) of the program does not change. When you do the same in a concurrent program, the global meaning of the program may change (the semantics is not "compositional" or "modular"). This is the right notion of a program that should be used in motor control. The issue is not discrete vs. continuous but sequential vs. concurrent/parallel. You may find it interesting reading the Turing award lecture of R. Milner, Comm. of the ACM, Vol. 36 No. 1, 1993.

Distributed control: The smaller is the spatio-temporal scale you employ, centralized become distributed. The "action potential of a neuron" is a word to describe a distributed phenomaon if you look close enough. There is no reason to believe that this does not go both ways and that single neuron is the only locus of control. (We had this thread before but I forgot the conclusion).

Re: Controlling for a process:

A: What are you doing?
B: (while eating, digesting, reading, writng, laughing, crying,
arguing, walking, running, sleeping, driving, typing..)

I'm dying.

Oded Maler

Date: Thu Feb 04, 1993 9:09 am PST
Subject: articles of interest

[From: Bruce Nevin (Thu 930204 11:19:22)]

New Scientist for 1/16/93 has an article on analog chips modelling neural currents in neurons, pp. 29-33.

And _NS_ for 1/9/93 has a short note on p. 17 about how female doves must coo during courtship to trigger release of eggs for fertilization. Disable the cooing, no eggs. Play recordings of cooing, she ovulates, especially if it is her own cooing that is on the recording. Doesn't matter if she's deaf, so the environmental feedback may be through sensors closer to the ovaries than to the brain.

Bruce bn@bbn.com

Date: Thu Feb 04, 1993 9:40 am PST
Subject: control article

[From Rick Marken (930204.0900)]

Hans Blom (930204) --

> William S. Levine, Gerald E. Loeb's paper 'The Neural Control of Limb Movement'

>' the study of biological systems should not be confined to testing
>whether their performance is compatible with control schemes invented
>to date but must include detailed examination of their inner workings
>to discover new types of control'.

What do they mean by 'new types of control'? There is only one type that I'm aware of -- maintenance of a perceptual variable at a specified reference level in the context of variable disturbances. Levine and Loeb give no evidence of understanding that it is a perceptual variable that is controlled (by whatever means -- bang/ bang, continuous output, lagged output, etc) and they seem to assume that they already know which biological variables are con- trolled; the only problem (they think) is to figure out the mechanism by which control is implemented. Levine and Loeb (didn't they have a run in with a young kid in Chicago some years back?) seem to have forgotten to mention step one in the study of biological control systems -- the test for the controlled variable. How can they compare the performance of a living system to know "control schemes" if they don't know what is the living system is controlling? Do they explain how they know what variables are controlled by the high jumper that they mention? If so, how do they know?

Sounds like another nice entry for the Devil's Bibliography.

Best Rick

Date: Thu Feb 04, 1993 10:34 am PST
Subject: knowing about social agreements

[From: Bruce Nevin (Thu 930204 11:44:20)]

If I don't respond to a note embedded in a longer posting it is because I am not getting to reading everything through.

(Bill Powers (930203.1000)) -- Bruce Nevin (930203.0849) --
>
>>Here's one aspect of the problem: the extent to which we live
>>in a world of the imagination, strongly colored by our
>>verbalizations, is I think insufficiently appreciated. The
>>grease of social agreements lets us get away with ignoring (or
>>approximating, by convention) a little raspiness in our
>>feedback through the physical environment.

>How do you, personally, know there are any social agreements?

I don't know in what direction that is a leading question for you, but for starters there is a fairly low-level social agreement behind the fact that you and I would both use the word "cat" while together scratching your cat (or one of mine) under the chin. We would agree that it was the same word (a repetition), and the same cat (identical, not a repetition), and I suppose those are social agreements too, but not the sort that we had in mind. For a great many of the perceptions that we as individuals associate with the category "cat", if we each used words that we individually associated with those perceptions, we would agree that those (words associated with) perceptions indeed pertain to (our respective perceptions of) cats, whether or not we and a cat were currently in one another's company. Social agreements thus outlive the circumstances that occasion them.

They long outlive them. We know well what a "bigwig" is, a century or more after the wealthy and their representatives wore big wigs in this country. We know what it is to "take the bull by the horns" many centuries after the worship of Mithra (the "soldiers' god") was spread with the Roman Empire (this idiom is found in literal, word-for-word equivalents in a dozen or more languages of Europe).

But we don't make these agreements, we use them, having learned them in childhood as the category perceptions used by people around us, which we had to learn to use in order to arrange cooperation with them.

The peculiar comparison expressed by "the more . . . the X-er" is conventional in English, but not in many other languages. We easily and without effort associate imagined perceptions with:

    The more people come, the merrier the group will be.

We even have no difficulty with the conventional reduction:

    The more the merrier.

If we have the perceptions that are associated with those phrases and wish some other person to have them too, we produce utterances like these flawlessly, and repeatedly find our confident expectation that they will be understood rewarded. By again and again relying upon such agreements and finding them dependable we continually reaffirm them in ourselves and (we presume) in others. This seems to be how we learned them in the first place.

We rarely notice them except, say, when one who is not a native speaker breaks from seemingly idiomatic English to stumble over an agreement such as this to which he is not party, that is, which has no corollary in his language. At this point, I will violate a social agreement that is very strong in this country of immigrants (except for the natives), and by way of illustration call attention to just this slip in a recent post. Oded, I hope you will pardon my transgression of good manners:

>(Oded Maler 930204) --
>
>Distributed control: The smaller is the spatio-temporal scale you
>employ, centralized become distributed.

Bill, if you would urge that these are not social agreements, or that we can't know that they are, please elaborate. I will respond as timely as I am able.
Bruce bn@bbn.com

Date: Thu Feb 04, 1993 10:41 am PST
Subject: Blunders^2; modulator;Mars trips and optimal control

[From Bill Powers (930204.0900)] Avery Andrews (930204.1224) --

"an orderly sequence of symbol strings" is, as you say, a lot of blunders rolled into one.

There are some blunders that are so extremely far off the track that one wonders whether the effort of correcting them is worth while. People who aren't prepared to admit ignorance aren't very likely to want an education. All that arguing against such persons does is drive them to wilder and wilder flights of fancy as they try to find a defense, not really caring much if their argument holds water and being much more concerned about winning or appearing to win. There's no doubt that such people can be very bright. That's the main problem. They've been used to understanding everything before the sentences are finished, because basically they're in an easy field of study and have far more mental horsepower than is required for it. When they come up against something new that can't be grasped instantly, they simply don't believe that their immediate understanding could be wrong. That's a new experience, and one they don't much like. Actually having to do some hard homework like one of those plods who gets Cs is unthinkable.

I'm an authority on such people, having been one of them.

>On the other hand, in the next section, they make what strikes
>me as a sensible argument that processes that may look like
>feedback control w.r.t. a set-point aren't necessarily so, and
>that respiration rate-stabilization in fact isn't.

Are you sure they're not just assuming that reference levels are constant for all control systems? If they believe as many biologists do, that "set point" means something set to a fixed value (homeostasis), they're arguing against a straw man, because control theory doesn't claim that reference signals are fixed. And if they think respiration rate is supposed to be controlled, they're talking about control of output.

For something like respiration rate, there's no a priori reason to suppose that it's under control. Automatic respiration is run by a oscillator of variable frequency and amplitude. These may or may not be controlled variables in themselves. For sure, however, this system is an output function used to control other variables, like blood CO2 and oxygen tension. On the other hand, when conscious control of respiration is involved there is plenty of evidence that the breathing muscles can be used to control inhaling and exhaling in any pattern you like, even one that overrides (for a while) the autonomic control systems.

>Then there's an interesting-looking argument by G. Hinton
>arguing that it is in fact a good idea to precompute torques,
>and discussing various features of muscles that make this
>easier to do than you would think (for example, the viscosity
>properties of the force-velocity relations give you what is in
>effect instant feedback control over the velocity of a limb:
>if the velocity is less than expected, the force exerted by the
>muscle will be higher, so the velocity lag will be corrected.

If someone will send me a copy of this article, I'll critique it.

RE: Picture proposal.

Wolfgang Zocher's SIMCON is almost ready for distribution. When it's done, you can try modeling that diagram to see what it actually would do when used to control an external variable.

My prediction is that with a leaky integrator in the input function, the system (if capable of good control) would make the external variable behave as an imperfect first derivative of the reference signal. Raising the reference signal from one level to another, for example, would cause the input quantity to increase sharply, then fall back almost to its original level.

The effect of signals entering the output "modulator" will depend on their effect on the output function. If they simply add, they will act as disturbances that the control system will oppose. If they act like multipliers, they will change the output sensitivity of the control system, changing its loop gain. Either effect, if not so large as to overwhelm the control system or reduce its loop gain below the necessary amount, will disturb the operation of the control system, but the perceptual signal will remain in a near match to the reference signal.

Beware of proposing block diagrams that you haven't actually tested as a working model. Closed-loop systems can fool you.

Hans Blom (930204) --

The fastest way to Mars that you propose assumes unlimited propellant. If the store of propellant is finite, the fastest way is to take off by using about half of the energy in a single impulse blast (say, out of a cannon) to put the spacecraft into a transfer orbit, and using the rest of the fuel at the last instant to bring the spacecraft to a halt just at the surface of Mars. In this way as little fuel as possible is used to combat gravity in taking off and landing (note that if the thrust is kept just equal to the weight of the rocket, all the fuel will be expended without lifting the rocket at all). Of course after this sort of trip you could then open a stopcock in the spacecraft and drain the crew out. The Saturn rockets had to be so large because there was a limit to the acceleration the living astronauts could stand.

Human control systems are pretty close to the design limits set by the materials used. It's possible, for example, for an arm muscle to pull itself loose from its attachments to the bones, if feedback is lost and an energetic movement is attempted. Even with an intact set of control systems, tendons and muscles can be ripped loose if an emergency situation results in sending abnormally large reference signals to the spinal motor neurons.

The "substantial deviation from the optimal control model" that Levine and Loeb mention may not be a deviation from what is optimal for the whole human system using the control system. Control models of an arm usually propose the application of torques at each joint, but in the human system there are no motors at the joints. Instead there are nonlinear muscles attached in clever ways that product many kinds of torques, some through clever linkages (as in the two bones of the forearm) and some by having the muscle wrap around the joint in a strange way (like pronator teres or the biceps).

Even the muscles work differently from the servo motors that engineers use. They don't apply forces directly, but by shortening the contractile elements in the muscle to alter the resting length of the series spring component. In principle, a movement could be carried out by suddenly shortening all the contractile elements in a muscle and storing energy in the spring components, and then letting the spring components execute most of the movement without any further expenditure of muscle energy until time for deceleration. Actual movements work somewhat in this way. This is something like the solution for maximum rocket efficiency given a finite fuel supply. In fact, the human system is far more efficient than any robot so far invented; it moves 100 to 200 pounds of weight around all day expending only 2 or 3 kilocalories of energy and using less than 0.1 horsepower of total muscle output power. And the fuel supply has to support not just the muscles, but the brain and the general metabolic requirements.

The reason a human being can't perform a mathematically optimal jump is simply the rocket problem: you would need to produce an impulse of muscle force of zero duration and infinite amplitude. That would hardly be a feasible solution for a servomechanism, either.

The "feedback too slow" argument turns up even here, doesn't it? Actually the speed of feedback in human control system is just right -- to explain the behavior we see.

>But stabilizing control is incompatible with top performance, such as
>in sports. In maximum height jumping, only the maximum height of the
>jump is important, not the full trajectory.

Human beings hardly ever control the "full trajectory." They control the variables that matter to them. Actually Rodney Brooks has the right idea here: don't plan trajectories, avoid obstacles. It isn't necessary to know where obstacles will be if the system has sensors that can detect proximity to an obstacle.

"Stabilizing control" is something of a misnomer, suggesting that all that a control system does is to keep something constant. More generally, it makes the perceptual signal track the reference signal. This means that a control system for producing a directed force (as in throwing a ball or launching a high-jump) can make the sensed acceleration have the right magnitude and direction right up to the moment of release. When we learn how these perceptions must change in order to have a desired result remotely or later, we vary the reference signals to repeat the experienced thrust as nearly as possible, and get pretty close. Of course if we got too close people would stop doing such things -- or they'd set the bar higher, or put the target farther away, until errors in control once again made the game interesting.

I think that when normal human movements such as walking are finally modeled fully, we will find that the system uses as little energy as possible, letting momentum and spring effects carry most of the movement through, with muscle contraction being used primarily to trim the result into a useful form. When we walk, we choose a pattern of walking to control that is as close to the zero-energy pattern as possible, given the higher-level goals of actually getting somewhere in a reasonable time. Only when we have some reason to get there faster, as in running a race, do the control systems try to produce patterns that cost a lot of energy. And even then, the patterns finally chosen are pretty efficient -- after all, the fuel supply and distribution have to suffice to get the body to the finish line.

Bruce Nevin (930204.0942) --

What do these people mean by "intentionality?" Is this more of that "directed toward" or "aboutness" stuff?

Best to all, Bill P.

Date: Thu Feb 04, 1993 10:56 am PST
Subject: Re: knowing about social agreements

[From Oded Maler 930204 - 18:55 European time]

[Bruce Nevin, few minutes ago]

*
* We rarely notice them except, say, when one who is not a native
* speaker breaks from seemingly idiomatic English to stumble over
* an agreement such as this to which he is not party, that is,
* which has no corollary in his language. At this point, I will
* violate a social agreement that is very strong in this country of
* immigrants (except for the natives), and by way of illustration
* call attention to just this slip in a recent post. Oded, I hope
* you will pardon my transgression of good manners:

I would, if only I knew what transgression is..

*
* >(Oded Maler 930204) --
* >
* >Distributed control: The smaller is the spatio-temporal scale you
* >employ, centralized become distributed.

I knew that this sentence is not correct but soemhow errors in English do not cause me any disturbance, I mean I feel that it is just a game, not the real thing. I don't care to wear T-shirts with stupid English or French slogans printed on them, and I'm sure I would not be wearing them if they were in Hebrew. If I don't focus intentionaly on Latin text, it is anyway just a meaningless sequence of symbols. An equivalent error in Hebrew would "shout" to me much stronger. (And here again comes the question of placement in the hierarchy, I'm sure my perceptual signal for Aleph is lower than that for 'a'.

Coming to the orignal point, the question is how these "social agreements" are arrived them, and how they are "realized" by individuals. How is the abstract aggregate entity called the English language with its rules and conventions is realized by the behavior of individuals.

--Oded

Date: Thu Feb 04, 1993 11:10 am PST
Subject: End-point control; the H-word; mindreading

From Greg Williams (930204) Bill Powers (930203.1000)

>End point control involves comparing what you're experiencing
>with what you want to experience and turning the error into an
>action that will make the error smaller.

But not necessarily "immediately" smaller -- you might have to go around an obstacle, but I still call that end-point control. As I envision end-point control, it can be very complicated (i.e., involve many possible sub-tasks), but it does NOT involve an error signal generated by comparing reference signal for "being on the path to" the OVERALL goal (overall task) with the current trajectory -- that is automatically taken care of at lower levels, with no need for asking the question (continuously) "am I on the right path?"

>Rate of change control, or transition control, is a little
>smarter. With this level of control, which exists in the crowd
>program, the organism can turn away from the configuration-goal
>in order to keep its velocity going, and this will lead it around
>most obstacles. Velocity control has no end point.

You have reified my more general notion of an end-point condition, which need not be spatial. I would say that end-point control for velocity tries to achieve the reference level currently set for the velocity, with no regard (at that level!) for errors in the trajectory involved in achieving that velocity ("end point"). I am just saying that I (still) don't see the need for computing an error measure for how close to (some ideal?) trajectory one is at any point. The error between where one is and where one wants to be (again, not just a spatial "where") should be enough, coupled with lower-level control circuits which adjust the actual trajectory in the light of disturbances.

>It seems essential that high-level errors be converted into
>changes in lower-level reference perceptions in a series of stages.

I think so, too.

Maybe we are saying the same thing differently. I am only saying that there appear to be long-term errors (hours, years) endemic to the higher level control circuits. One wants to be in Paris; one is in Colorado. Ergo, Paris is (with current technology) several hours away, and the error signal will be non-zero for at least that time. But you seemed, in replying to Oded, to be denying that "overall" error and its persistence for some time, by speaking of "what one can control" from moment to moment. I don't deny that path control is occurring at the lower levels. Do you think end-point control (as I've defined it above) is NOT "driving" the lower levels?

>I've been trying to explain here why I think they're necessary.
>But beyond the logical problem of converting an abstract error
>into more detailed goal-adjustments, there's another fact that's
>just as important as the abstract principles involved. It's that
>people actually do control for the form of a process.

Yes, they do. Do they not also control for the COMPLETION (satisfaction, terminus, whatever) of such processes? It appears to me that they do, and it further appears to me that such completions don't occur (in general) instantly. And the error at that level remains non-zero for a period of time.

>>... in general, higher- level goals are NOT satisfied, but only
>>in the process of BEING satisfied. Why do you want to maintain
>>that goals are practically always satisfied over time?

>I didn't say they ARE always satisfied over time. I just said
>that a wise person will select higher-order goals so they can be
>achieved immediately: don't aim to be a millionaire; aim to do
>all the things that accumulate money rapidly.

Well, I THOUGHT I heard you say more than once on the net (once fairly recently) words to the effect that organisms typically have essentially zero errors throughout their control apparatus.

>If you adopt the time-scale appropriate to the speed with which
>any level of control can correct its errors, the process of error
>correction occupies the shortest distinguishable time unit and
>can be ignored. The kind of process control I am talking about is
>concerned with variables that change very slowly in comparison
>with the time it takes to correct an error in the process. I'm
>talking about controlling perceptions, where the perception
>itself represents the presence of an ongoing process.

That's what I wanted to hear. I have no problems with that clarification.

>>You misunderstand. I haven't seen sufficient evidence to
>>justify YOUR vision of the way peoples' control structures

>>operate as either the ACTUAL way or the only POSSIBLE way.

>In the binary choice between sufficient evidence and insufficient
>evidence, I suppose you're right. I think I have some reasonable
>ideas, and many informal examples that seem to fit them. But that
>falls between "right" and "not right" so I suppose it's in
>territory that you don't acknowledge as meaningful.

I agree that you have (MANY) reasonable ideas and informal examples which support them. Of course weight-of-evidence counts. There is no such thing as empirical certainty, so ALL decisions depend on the gray territory. That territory is certainly meaningful. But brain modeling and neurophysiology are both young sciences, and I think that, while it isn't too soon to make inspired guesses about the details of brain organization, we are quite a ways from critical experiments to choose among those guesses. That might be a frustrating point of view to you, but the alternative is leaping to a very possibly wrong conclusion and then being absolutely devastated if it turns out wrong. High-stakes gambling never appealed to me. Instead, I appreciate Chamberlain's method of multiple hypotheses as a way of hedging one's bets. Then, one might never win big, but one will never lose big either. The Powersian H in PCT is still in my hand -- and I expect to keep it there for a long time. But I think it prudent to draw a few other cards, too -- and to be able to discard any of the cards when appropriate.

Yet I suspect I would feel differently if HPCT were my life's work.

>While adaptive control and reorganization and gain control and
>other subjects are important, I think it is premature to worry
>about them. I'm still trying to get people to see the
>significance of ordinary control, the simplest kind.

I think that is a good strategy. I also think your strategy of suggesting that others should attempt to concoct other sorts of H models for PCT is excellent.

>Even in our group, everyone has some idea they're not willing to look at very
>closely -- the ideas that are "obviously right." So it goes.

Yes, everyone. Mea culpa.

>Avery.Andrews 930204.0940

>>So why not phone him? The paper affords a perfect entre for direct
>>communication. You could ask him what he thinks of your ideas and what

>Because if his articles are obscure, what he says in a phone-call from
>out of the blue will be even more so, & the entire experience is
>likely to be extremely embarrassing. If the writing looks obscure,
>you just publish something claiming not to understand what it means, &
>leave it to the author to clarify, if they are able and willing.

>>How do you know until you've tried? Very few attempts to communicate
>>with "devils" have been made to date by PCTers.

>Because I am one myself and I therefore know what they're like. I think
>a number of conclusions can be drawn from (a) typical referee responses

>to PCT papers (b) the way the communications with Randy Beer and
>David Chapman proceeded. I see no reason to think that unsolicited
>phone calls would turn out any better.

The proof of the mindreader is, of course, in the accuracy of predictions. Perhaps no devil-critique paper even needs to be written, if some PCTers can accurately predict the devils' replies to it?!?! But the accuracy of those predictions has yet to be tested sufficiently, in my experience.

As ever, Greg

Date: Thu Feb 04, 1993 11:38 am PST
Subject: rt stack, finger poppin'

[From Rick Marken (930204.1000)]

The reaction time Hypercard stack is comin' along. The main discovery to date is just a confirmation of what Bill P. said some time ago which is basically that lags and system response time act as a low pass filter. The stack is set up as a two level system; level two controls perceived position by varying the reference to the level 1 system which controls perceived velocity by producing a force output; so it's a model of a "finger moving" system. I can vary the transport lag (the time it takes for the input to be "transported to the output) and the system response time (the rate at which one signal (error) is transformed into another (output)). I can apply a step or impulse disturbance to the controlled variable (the position of the "finger") and I can vary the rate of change of the position reference signal (a discrete change; the reference signal value is instantly changed from one value -- say, 20 -- to another -- say, 40).

The rate at which the reference signal changes determines the amplitude of the (now smoothed) changes in finger pos- ition. The faster the changes in the reference, the lower the amplitude of the position oscillations -- the higher frequencies are being filtered out. Both lag and system res- ponse time influence this filtering -- and they have slightly different effects on the "shape" of the temporal variations in finger position. I have not yet checked on the effects of a disturbance while reference signal changes are occuring.

I have verified that the low pass filtering seen in the model also happens in a person (me). Bill mentioned a similar demo with the arm but I did it with a finger so you can do it too, right there at your computer. Put your left hand up with thumb and middle finger in a cresent with the ends about 4 inches apart. This is the target separation. Now move the index finger (only) of your right hand back and forth between these two point at the rate of about 2 swings per second. No problem, right? Now gradually speed up the oscillations of the index finger -- trying to keep the index finger moving to each of the two targets. What you will find is that, as you speed up the oscillations, the amplitude of the oscillations decreases. So you can change the position refer- ence very quickly but when you try to change it too quickly the finger position will not end up where it is "intended" to be.

Actually, now that I think of it, this can also be done with a mouse moving a cursor between two points. Could this be a way to get at the lag and system response time of the arm movement system? Bill P. what do you think?

By the way; I've found a clever way to speed up my HyperCard scripts; run 'em on a Quadra.

Best Rick

Date: Thu Feb 04, 1993 12:03 pm PST
Subject: unintended intentionality

[From: Bruce Nevin (Thu 930204 14:18:42)] (Bill Powers (930204.0900)) --
>---
>Bruce Nevin (930204.0942) --
>
>What do these people mean by "intentionality?" Is this more of
>that "directed toward" or "aboutness" stuff?
>--

Beats me. I don't remember posting anything about intentionality, and can't find any 930204.942 in my mailbox from any source. The only occurrence of "intentionality" turned up by grep is your (930203.1000) reply to Avery. Am I missing something?

Bruce

Date: Thu Feb 04, 1993 2:19 pm PST
Subject: TO ALL "BILLs" - RKC

Bill Silvert (930201)

ADDRESSING MAIL
Thanks for your suggestion about addressing mail to a principal recipient. It was thoughtless of me to assume "Bill" would be sufficient. I hope that you and any other "Bill's" will accept my apology.

Regards, Bob Clark, "RKC"

Date: Thu Feb 04, 1993 2:45 pm PST
Subject: Intentionality

[From Chris Malcolm]

Bill Powers writes:

> What do these people mean by "intentionality?" Is this more of
> that "directed toward" or "aboutness" stuff?

Yes.

I just happen to have written (ten minutes ago!) some notes for first-year students introducing "intentionality", which I will append. First I'd like to explain why I think intentionality is an important problem in the architecture of creatures both biological and artificial, and what part PCT plays in this.

A sore philosophical problem in AI and robotics is how to provide a robot (or artificially intelligent system) with intentionality, in the sense of having inherent semantics of its own, as opposed to the borrowed semantics provided when a human observer interprets something as having meaning. For example, Searle's "Chinese Room" argument is concerned to deny the possibility of inherent semantics to any computer system, ever (by virtue of running some program), which many have generalised into a denial that artificial intelligence (in the sense of giving machines ideas) is possible.

I'm developing a paper which shows how to get intentionality into robots by means of two kinds of machinery, the first of which consists of a collection of hierarchical control systems. I argue that describing the behaviour of a device or creature containing goal-seeking mechanisms in teleological terms is not only useful (as in Dennett's "intentional stance") but the _only_ way of offering a principled scientific description of the behaviour.

I then argue that the goals of the goal-seeking mechanisms offer the most natural set of terms in which the behaviour of the creature can be organised, in other words, if the creature is to be provided with some capacity to reason about its perceptions and actions in the world, as part of deciding what to do, then these goals provide the most natural high level terms in which to categorise and articulate its perceptions and actions. These terms are also naturally grounded (in Harnad's sense of "symbol grounding") by the mechanisms from which they are derived. Of particular importance is the way that this avoids the problem which I think unseats Harnad (and most roboticists, with the honorable exception of Brooks and some others) -- thinking of an architecture with two hierarchies, one dealing with input (perception) and the other output (action), in which lower-level cross-connections between the two are optimising adjuvants, rather than -- as I think -- the crucial feature of the architecture.

Thus by a combination of reasoning implemented on top of goal-seeking machinery the robot is equipped with intentionality. In this connection I think it will prove particularly useful to describe the goal-seeking machinery in PCT terms: as controlling perceptions.

Now follows the bluffer's guide to intentionality:

"Intentional" is a technical philosophical term related to "intention", but not derived from it in meaning. It refers to the pointing quality of symbols, the fact that a symbol means something as well as simply being something (like a mark on paper or an arrangement of transistor states). The "aboutness" of symbols is their intentionality.

The philosopher Brentano (1838-1917) suggested, in what is now known as "Brentano's Hypothesis", that "intentionality is the ineliminable mark of the mental". Here he meant more than simply the "aboutness" of meaning -- he meant the conscious act of meaning. However, in modern philosophical usage it means simply "aboutness", meaning, or semantics as opposed to syntax, without necessarily presuming consciousness. Quite what it consists of is a thorny and much disputed philosophical problem.

Scientists have always been very wary of anthropomorphising or being subjective when describing things. So it is usually held to be wrong, when describing the behaviour of an animal, to say that it "thinks" something, "believes" something, or "intends" something. Nevertheless, using that kind of language results in very compact and powerful descriptions of goal-directed behaviour. For example, a chess-playing machine can be said to be trying to gain control of the centre of the board, while being careful to defend against your possible queen-side knight attack, as a predictively useful description of its behaviour. If someone objected to this language on the grounds that machines can't "try" or "be careful" etc., trying to adopt a lower-level description of the same behaviour -- say in terms of the machine's main software procedures and data structures -- would be both much longer and much more difficult to understand.

Therefore the modern philosopher Dennett has devised what he calls the "intentional stance". This consists of putting to one side, for the moment, questions of subjectivity and anthropomorphism, and using directly intentional and purposive language to describe the behaviour of complex machines (or simple animals), "as if" they really did have proper intentions, plans, beliefs, and so on. The justification for this is that it provides a compact and powerful high level of of description of behaviour without specifying at all how the behaviour is implemented.

It is possible to describe computer programs at a variety of levels. The lowest is the code. A higher level description is in terms of the data structures and what is to be done with them. The highest level is the specification of what the program is meant to do, without any reference to algorithms or data structures. This could be said to be the intentional level of description of a program.

Chris Malcolm

Date: Thu Feb 04, 1993 3:08 pm PST
Subject: Language conventions; process control

[From Bill Powers (930204.1400)] Bruce Nevin (930204) --

My question about how you know what the social conventions are elicited your justifications for what you believe about them. Oded Maler picked up on the intended meaning of my question:

>Coming to the original point, the question is how these
>"social agreements" are arrived [at], and how they are
>"realized" by individuals. How is the abstract aggregate entity
>called the English language with its rules and conventions
>[] realized by the behavior of individuals. (Maler 930204.1855 ET).

You as a linguist have a special position with respect to beliefs about the social conventions. You cast them in terms of Harris's theory, saying that they entail operator-argument dependencies, zeroings, and so forth. These constitute the principles you see at work when you perceive discourse, and all these principles taken together constitute your social system concept of language. That system concept hangs together for you and it creates for you the concept of language as a coherent entity.

Avery Andrews, listening to the same discourse, perceives different principles at work, and for Avery, these principles taken together constitute his social system concept of language. Clearly, the language-entity he perceives and believes to be a correct perception of the social conventions is different from yours.

Mine is still different, because I begin with PCT concepts and hardly any knowledge of language in the manner of a professional linguist. I can only apply my principles to language as I experience it, and what I end up with as a system concept of the language entity is different from both yours and Avery's. I perceive the social conventions in a way consistent with my theory.

So what, then, is "the" social convention about language? You will find different conventions in different parts of the English-speaking world. Most of these people being non-linguists, their perception of what the conventions are must vary widely, many of them being incompatible with yours, Avery's, mine, and each other's.

The only thing we can say for sure is that each of these understandings of the social conventions about language is pragmatically sufficient for the task of communication. That is essentially a paraphrase of Wittgenstein, isn't it? However a person may perceive the systematic structure, the principles, and the rules of language, those perceptions have been honed over time until they adequately, for that person, explain what the person hears, and results in production of utterances that others consider meaningful and adequately structured. In effect, each person has a theory about the structure of language, but it is not necessary that any two of these theories be identical. All that is required for the successful use of language is that utterances constructed according to one person's way of perceiving rules, principles, and concepts of language fit into the listener's way of perceiving at the same levels of experience.

A linguist may come up with an understanding that is based on far more critical analysis of utterances than a layman would be able to construct (or nonverbally learn). This means that when the linguist speaks or listens, there is a much richer experience going on, with far more details and a far more self-consistent set of principles than the layman has, and a much more coherent sense of the whole phenomenon of language. But this is the linguist ordering the linguist's experiences to the linguist's satisfaction.

When a layman speaks or listens, what is heard is much less detailed and coherent, and the language shows it. After all, language is produced in the form of controlled perceptions, and a layman can control utterances only in the terms that the layman can perceive. The linguist's rich and consistent perceptions of language do the layman no good. The layman will speak in ways consistent with the rules, principles, and concepts that the layman has managed to work out by experience or thought, in process of making sense of what others say and getting others to demonstrate comprehension of the intended meanings. This doesn't mean that the layman could express this theoretical structure; it means only that the structure exists and is effective.

I don't believe that we can talk about "the" structure of language, or "the" social conventions of language. Each person has made sense of the structure and the conventions in terms of a privately-developed set of rules, principles, and concepts. All we can say about the commonality of these rules and so on is that the test of conveyance of meaning and acceptance of forms as being reasonably "correct" is not failed. This is not an indication that there are any social conventions known in the same way to all people. Each person thinks he or she understands the conventions correctly. Each person follows conventions that are a little or a lot different from what the person thinks they are. But no person, not even a linguist or a control theorist, knows how another person understands them.

So that's how I answer my own question.

Greg Williams (930204) --

>I THOUGHT I heard you say more than once on the net (once
>fairly recently) words to the effect that organisms typically
>have essentially zero errors throughout their control apparatus.

I sometimes leave out conditions on my statements, and so claim more than I mean to claim. I claim that in a hierarchy of control systems that is operating in some sense optimally, errors remain close to zero in terms of the duration of the specious present appropriate to the level under discussion. This does not apply, of course, to control systems that are not very skillful, or that are in conflict with other control systems, or that are presented with disturbances too large to resist, or that are reorganizing.

This leaves unexplained control processes that apparently take a very long time, like getting a PhD. Such control processes, by my definition, are operating pretty far from the optimum performance. And this is really true: how much control over PhDness does one have during all those years of school? Can we even speak of control of something that is a unique one-time event? If you fail to get the PhD, how do you correct the error? There are few ways, and I doubt that many people try them. The transport lag of the control loop is so long that one could hardly go around it three or four times; at best one can write the thesis again or take a course over, but not very many times. If the problem was with your high-school education, you're not likely to traverse that loop again.

Such long-term control loops do seem to exist in people, and I'm not entirely sure how to think about them. We could say that yes, these control systems exist, and no, they're not very good control systems, and can't be by their nature. But my leaning is toward a different conclusion. It's that we would be better off acquiring good control systems.

It's possible to set a goal of learning the subject-matter of a given course, today, this week, as well as one can possibly learn it. Goals can be set and achieved for amount of study, for methods of study, for choice of things to study that are interesting and likely to be learned. These immediately achievable goals can be consistent with all of a person's other goals, so that the whole picture of one's life in school is satisfactory on all counts. If that's true, then why worry about getting a PhD? If you go on learning, you will acquire the knowledge and skill that you want whether or not the PHd comes along somewhere in the middle of the process. If it doesn't come along, then getting it probably required you to deviate from your learning process. And the chances are that the PhD will take care of itself if you do all the rest. If by some misfortune you are not really interested in the learning itself, but just want the PhD for practical reasons, your life while getting it is going to be hell. Why go through that?

I guess what I'm saying is that if you organize your goals, and state them, in such a way that they can all be satisfied more or less immediately, the final result will be better, and your life will be better, than if you set up shining but distant goals and experience your life on the way to achieving them as one long huge error signal. Protracted large errors are a sign of a bad control system.

I guess we have reached some agreement on process control.

>... while it isn't too soon to make inspired guesses about the
>details of brain organization, we are quite a ways from
>critical experiments to choose among those guesses.

While we're all guessing, I think it's best to make our guesses as informed and coherent as we can, taking into account as much of experience as we can handle and trying not to contradict basic knowledge that is reasonably well established.

>I appreciate Chamberlain's method of multiple hypotheses as a
>way of hedging one's bets.

I'm always considering multiple hypotheses where I can think of any. For the basic phenomena of control, I can't think of any. All the other proposals I've heard about have something wrong with them -- they ignore facts, they don't actually say anything, they predict things that don't happen, they're too imprecise to be tested, and so forth. I wouldn't adopt multiple hypotheses just to have more than one. They have to be good ones.

Bruce Nevin (930204.1418) --

>I don't remember posting anything about
>intentionality, and can't find any 930204.942 in my mailbox
>from any source. The only occurrence of "intentionality"
>turned up by grep is your (930203.1000) reply to Avery. Am I
>missing something?

Try this:

===

[From: Bruce Nevin (Thu 930204 09:42:00)]

Penni, if you're still here, this looks like a natural for you:

Linguist List: Vol-4-66. Wed 03 Feb 1993. Lines: 125

Messages---------------------------------------

Date: Wed, 3 Feb 93 14:21:31 EST
From: rambow@unagi.cis.upenn.edu (Owen Rambow)
Subject: Workshop: Intentionality and Structure in Discourse
Relations

CALL FOR PAPERS

INTENTIONALITY AND STRUCTURE IN DISCOURSE RELATIONS

Etc.
===

Best to all, Bill P.

Date: Thu Feb 04, 1993 3:34 pm PST
Subject: telephoning devils

[Avery Andrews 930205.0938] (Greg Williams 930204)

Well, my internal model of Real Professors tells me that publications are much more likely to have an effect than phone calls (e.g. some chance, rather than no chance at all). Partly because the target has some time to think about what has been said, and come up with a genuinely useful response (I think this is much more time-consuming than the conventional rules for spoken debate allow for), partly because they have a strong motive to respond (so as not to be humiliated in public), and partly because you then get a chance to make an impression on the milling crowd of uncommitteds, who are the guys you're supposed to win over in order to pull off a scientific revolution.

Avery.Andrews@anu.edu.au

Date: Thu Feb 04, 1993 3:46 pm PST
Subject: Intentionality

[From Rick Marken (930204.1430)] Chris Malcolm --

>I just happen to have written (ten minutes ago!) some notes for
>first-year students introducing "intentionality"

This is a timely topic. I am planning to attend the Claremont College Conference on Consciousness and Cognition (March 5th). I was invited by Bill Banks who was the editor of my "Hierarchical Behavior of Perception" paper (which was rejected with encouragement to rewrite).The only talk I want to hear is "Why volition is a foundation issue" by Bernard Baars. Volition and intention seem like synonyms to me. So the following statement by Chris puzzles me:

>A sore philosophical problem in AI and robotics is how to provide a
>robot (or artificially intelligent system) with intentionality, in the
>sense of having inherent semantics of its own, as opposed to the
>borrowed semantics provided when a human observer interprets something
>as having meaning.

If Baars starts talking like this about volition then I'm in real trouble (understanding-wise). What does intentionality have to do with "inherent semantics"? I think of an intention as a want, desire or purpose. Intentionality (to me) refers to the behavior of systems that have wants, desires and purposes. You can tell the difference between behaviors generated by intentional and unintentional systems using the test for the controlled variable. Is this a different kind of intentionality than the AI/robotics kind?

>I'm developing a paper which shows how to get intentionality into robots
>by means of two kinds of machinery, the first of which consists of a
>collection of hierarchical control systems.

So PCT puts "inherent semantics" into a system? How?

>Thus by a combination of reasoning implemented on top of goal-seeking
>machinery the robot is equipped with intentionality. In this connection
>I think it will prove particularly useful to describe the goal-seeking
>machinery in PCT terms: as controlling perceptions.

This sounds good but I'm not sure I understand it. I'd like to understand intentionality (and volition, too) from an AI/Robotics perspective in order to prepare myself a bit for the Baars talk (which, I bet, is likely to come from that direction).

Best Rick

Date: Thu Feb 04, 1993 4:48 pm PST
Subject: waggling, intentionality, Kugler et al

[Avery Andrewsx 930204.1104] (Rick Marken 930203)

I have a program that waggles a finger too (currently Turbo C, but supposedly written so I can port it to Unix-XWindows without too much drama), and while it basically does what Bill says it ought to, it's very unstable at the *beginning* of the waggle, even for rates of 6wps, which I can attain, with lags on the order of 60ms (round trip) which seems reasonable. I've tried to be semi-serious about taking account of force-velocity relations for muscles, as well as time-constants, but of course may well have screwed things up in any number of ways. In particular I *don't* control for low velocity when near the target, but just derive a torque from the error-signal. So maybe things will improve when the control-circuits are more realistic (e.g. armdemo-like).

As for `intentionality', my one-sentence definition would be that it is whatever is fundamentally mysterious about how (real, human) symbols work. But Chris Malcolm's proposal certainly sounds interesting, tho I'm not sure off the top of my head if it does everything that somebody like, say John Searle expects out of the notion.

) Oded Maler 930204)

>The only linguist paper I've read was Gross' "on the failure of
>generative grammer"

This paper was basically obsolete when it was written (early-mid seventies), and much more so now. Its criticisms apply to the `classic' Aspects-of-the-Theory-of-Syntax type model with considerable force, but not at all to the more recent `lexically-based' ones. People who want to get a more contemporary impression of what the field is like should perhaps look at Robert Borsley's textbook _Syntactic Theory_ (Blackwell's, I think), and the Shopen (ed) volumes _Language Typology and Syntactic Description_. The Borsley book presents two current approaches to syntactic theory (with complementary strengths and weaknesses), the Shopen volumes take an informal but theoretically-based look at language-typology.

What kind of impression did Kugler make on the rest of the audience? My impression of them is that what they actually do is good and interesting, but what they say about other people is often monstrously bad and irresponsible. I suspect that they would be very vulnerable to a forcefully pursued, common-sense-based attack.

Avery.Andrews@anu.edu.au

Date: Fri Feb 05, 1993 10:22 am PST
Subject: Re: waggling, intentionality, Kugler et al

[From Oded Maler 930205 13:55 ET] Avery Andrews 930204.1104

* >The only linguist paper I've read was Gross' "on the failure of
* >generative grammer"
*
* This paper was basically obsolete when it was written (early-mid
* seventies), and much more so now. Its criticisms apply to the
* `classic' Aspects-of-the-Theory-of-Syntax type model with considerable
* force, but not at all to the more recent `lexically-based' ones.

Just an impression: when you put on your professional linguist hat, you sound more like the motor-control professionals you so nicely quote and analyze. On the surface it looks like someone responding to your feed-back too slow article by saying that you ignore a decade of development in chaotic doubly-periodic attractors for dissipative systems whose Lyapunov exponents can be embedded in Lie groups of prime order etc. (This is a parody. Maybe you are right in what you say, and maybe this hypothetical mathematical crticizer is correct in his critics of PCT).

But still, being already there, in what sense was it "basically obsolete" at 78-79? Was it then commonly recognized that the theory of generative grammer is completely inadequate to explain real language phenomena (as the paper claims)? If so, where did Chomsky or any of his followers state it explicitly? (The latest reference to Chomsky is, btw, 'Reflections on Language' (1975)). Or has the idea of generative grammer dissolved quietly and has been replaced by a radically different paradigm? I really don't expect a comprehensive explanation of the difference between the old models and "recent `lexically-based'" models, and why the latter are immune to the previous criticism, but a short historical sketch of your view of the paradigmatic dynamics might clarify the semantics of your claim. But anyway it is not necessary, we may all pass on the inherently-unspeakable with (feed-back coltrolled) silent hand waving :-)

* What kind of impression did Kugler make on the rest of the audience?

... (I used to make lunch-time polls almost every day. I don't think everything should be said in public). But in total, I found him more thought-provoking than many other experimental and theoretical psychologists of various sorts. I think

I can find his e-mail somewhere and maybe you can ask him for his opinion about your "misunderstandings".

(oof, since Bruce's comment yesterday I edit my comments and lose some of my spontanity (sp.?))

Apropos the wild speculation the grammer descends from motor programs (in which I believe) I have a wilder one (even less well-defined) : "It's all metabolism", it all lies in the control of this periodic process.

--Oded

Date: Fri Feb 05, 1993 11:16 am PST
Subject: 93rd Birthday; Suitcase Flinging

[from Gary Cziko 930205.1434 GMT]

Bill Powers (930203.1000) informs us:

>Starting on Sunday, Feb. 7, I'll be gone for 10 days: a trip to
>Green Valley AZ for by father's 93rd birthday on the 8th. . .

This is great news. It gives us hope that you will also be with us for a very long time to come. It would be fun to hear of all the things your Dad had done and does which "statistically" is not conducive to a long life. Hopefully he is a beer-drinking, lard-eating, cigar-smoking couch potato.--Gary

>Martin Taylor (930202 13:45) said:

>>But "flinging it into the air" is a common result of picking up something
>>you thought to be heavy, but isn't.

Rick Marken (930202.1200) responded:

>I have NEVER had this experience with a suitcase or any other
>lighter than expected object. But thanks for the tip; I'll
>watch out if I'm ever checking in a very large, very light
>suitcase in Canada.

Rick, you obviously have not travelled into airports where there is a very high rate of luggage theft. I've arrived at some airports in Africa where about the half the bags are emptied before they reach the luggage claim area. Here the ceiling is pock-marked with craters from travellers flinging their suitcases in the air.

Seriously, I think Oded Maler interpretation is on-target. We adopt certain postures in anticipation of the amount of force needed to do something and we can be wrong and have to try again. This happened to me the other day as I left my office building. I pushed on the outside door to open it and nothing happened. I then had to change my posture, lean into the door at a steeper angle and push harder. That worked as I pushed into a very strong steady wind. Had I used the steeper angle when there was no wind, I probably would have fallen flat on my face as the door would swung open too easily.

We certainly have expectations of how much perceived force we will need to accomplish things and adopte appropriate postures (body configurations), but I don't see how this argues for preplanned outputs.--Gary

Date: Fri Feb 05, 1993 11:56 am PST
Subject: knowing about social agreements

[From: Bruce Nevin (Fri 93025 10:41:18)] (Maler 930204.1855 ET)

>Coming to the original point, the question is how these
>"social agreements" are arrived [at], and how they are
>"realized" by individuals. How is the abstract aggregate entity
>called the English language with its rules and conventions
>[] realized by the behavior of individuals.

(Bruce Nevin (Thu 930204 11:44:20)) --

>But we don't make these agreements, we use them, having learned
>them in childhood as the category perceptions used by people
>around us, which we had to learn to use in order to arrange
>cooperation with them.

>By again and again relying upon such agreements and finding them
>dependable we continually reaffirm them in ourselves and (we
>presume) in others. This seems to be how we learned them in the
>first place.

Also refer to Bruner's account of how children are taught control of language (and by means of language), which we have discussed in the past. (Jerome S. Bruner, _Child's Talk_.)

(Bill Powers (930204.1400)) --

What I said regarding social conventionality of language made no mention of operators and arguments, and only used the word "reduction" once in what I though was a pretty innocuous way, i.e.:

>We easily and without effort associate imagined perceptions with:
> The more people come, the merrier the group will be.
>We even have no difficulty with the conventional reduction:
> The more the merrier.

(Substitute "paraphrase" here if you like. Or omit it, it is not
critical for the point that was being made.)

I did this on purpose. So please, you do likewise, and leave out reference to technical matters of linguistics while we work on the question you raised, namely what social agreements are, how we arrive at them, and how we learn them when they are faits accompli upon which those around us depend, and upon which we must learn to depend in like manner if we are to cooperate with those around us.

Oded's question (quoted above) is, you say, a restatement of yours. Then my responses (also quoted above) are a partial answer to your question as well as his. They are not justifications for my theoretical positions. They are (sketchy) descriptions of the process, as I understand it, by which we as children learned the ropes -- ropes of convention, of mutually attuned expectations, that we now as adults work, seemingly without effort or thought, to our purposes.

When you do respond to what I have said, instead of to what you want me to have said, please bear in mind the distinction between the processes of learning social agreements as a child, the processes of using established (institutionalized) social agreements as means for effecting one's active control for social coordination, and the social agreements themselves as something pre-existing that the child must learn.

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

Re the cross-post from the Linguist Digest about the conference on intentionality: I didn't keep that. I don't know what they mean by intentionality, but the "intentional stance" stuff presumably provides the relevant buzz-words to justify a conference. When I cross-post notices about conferences, etc. as being of possible interest to others in CSG-L, I may have only a superficial understanding of what they will talk about in the conferences, etc. This is the case, for example, regarding the material from Al Boulanger that Martin liked so well.

Bruce bn@bbn.com

Date: Fri Feb 05, 1993 12:13 pm PST
Subject: waggling, intentionality

[From Rick Marken (930205.0800)] Avery Andrews (930204.1104)

>I have a program that waggles a finger too

It sounds more realistic than mine; for example, I don't produce a torque; just a linear force. The times I use are also probably not realistic -- but easily varied.

>As for `intentionality', my one-sentence definition would be that it is
>whatever is fundamentally mysterious about how (real, human) symbols work.

Well, I turned to Webster's New Collegiate Dictionary and found that "intentional" is must be one of those words that Humpty Dumpty paid double.

The main meaning of intentional is "1. done by intention or design". And looking for the definition of "intention" we find "1. determination to act in a certain way" -- which seems fine to me, although Webster obviously did not understand control theory; if he did he'd have used the word "perceive" instead of "act".

But Humpty seems to have paid for a second (and, for me, not obviously related) meaning of intentional -- more like the one that Chris and Avery seem to have in mind:"2 a. of or relating to epistemological intention" (Humpty must have paid a bundle for that one). If that definition seems a bit misty there is this one to help out "b. having external reference".

The definition of intentional ends with: "syn. see VOLUNTARY" So I was not that far off in assuming that intentional and voluntary are synonymous.

Now, my problem is to try to understand what the first meaning of intentionality (where it means "volitional action") has to do with the second (where it means "having external reference" or, "having to do with what is, in Avery's words,'mysterious about how (real, human) symbols work'").

My wife (an english major) suggested that we do sometimes speak about the "intent" of a particular sign or symbol. Perhaps I could understand this better if I read more Jung.

Best Rick

Date: Fri Feb 05, 1993 12:26 pm PST
Subject: Re: Intentionality

[From Oded Maler 930205 17:40 ET] Chris Malcolm

I have no intention to start a comp.ai.phil.-style of discussion but I can't resist trying to reformulate your suggestion.

What is the semantics/denotation/aboutness of a reference signal (inside an individual)? It is the set of states-of-affairs in the world, such that when sensed by the individual will cause a zero error in the corresponding comparator.

Still a problem with levels - should you attribute intentionality to a signal in a motor-neuron, or a thermostat?

--Oded

Date: Fri Feb 05, 1993 2:20 pm PST
Subject: Suitcase Flinging

[From Rick Marken (930205)] Gary Cziko (930205.1434 GMT) --

>Rick, you obviously have not travelled into airports where there is a very
>high rate of luggage theft. I've arrived at some airports in Africa where
>about the half the bags are emptied before they reach the luggage claim
>area. Here the ceiling is pock-marked with craters from travellers
>flinging their suitcases in the air.

You're right. My experience of such places is limited. However, I do have the experience of sitting in an office in which the floor is pock-marked with the craters from me falling off my chair laughing after reading your posts.

>Seriously, I think Oded Maler interpretation is on-target.

I agree, though I hate to because Oded said he believes in motor programs (you must foreswear such heresies in order to enter the realm of TRUE PCT). I have had the experience of lifting a suitcase too far (because I started to control for a much heavier suitcase than expected); but my higher level systems have always been able to change the force reference quickly enough to keep the suitcase from denting the ceiling; and we have higher ceilings here than in Africa (who ever goes to Africa??? Were you in the Peace Corps or something?)

Best Rick

Date: Fri Feb 05, 1993 2:22 pm PST
Subject: BALLISTICS & PCT - RKC

BALLISTICS
The subject of Ballistics was recently raised in connection with "coin snapping," shuffleboard, and Muhammad Ali. A "Ballistic" event consists of rapid acceleration of an object to a sizable speed followed by its simple, uncontrolled, coasting until it is stopped by its surroundings. (Impact on a target, air friction, etc.) A simple device is used in physics classes for demonstration.

SLOW FEEDBACK
Firing long range weapons provides an example of the importance of the "SPEED of FEEDBACK" the time required for system operation.

When guns were first used on targets that were beyond visual range, results were poor. Soon "Spotters" were introduced to report the results. Thus the gun became more accurate. This combination can be regarded as a negative feedback control system, even though the return signal is relatively slow compared to the speed of the projectile. It does not permit control of each shot, but provides improved control of the over-all performance of the gun.

Without Spotters, feedback was slow indeed, hours to days were needed to get reports. Adding Spotters reduced the delay, providing much faster feedback. Self-guided weapons are now available: cruise missiles, smart bombs. These work better yet, with much faster corrections. With these capabilities, they correct for aiming errors, possible movement of the target and varying winds.

TIME SCALE & "CLOSED LOOP" VERSUS "OPEN LOOP"
Analysis is influenced by the time scale selected. When times of the order of seconds are of interest (approximately the time needed for the projectile to arrive), there is no control without self-guidance. Here open loop analysis applies. Events are followed around the loop without treating the system as a whole. When events are examined in terms of the time for firing the gun several times (several minutes), closed loop analysis applies to each firing of the gun as the assigned target is followed. Assuming the necessary components are present, either closed or open loop analysis may be suitable according to the time scale of interest.

CONTROLLED VARIABLE
A primary question for any control system is: "What is the perceptual variable being controlled?" In this case it is the point of impact of the projectile. This variable is a combination of several perceptual variables used to specify location in terms that can be communicated to the gun crew. The observer's conscious attention is required in combining and communicating this information.

GENERAL EXAMPLE

This system, assembled for the purpose of controlling the impact of a projectile, can be used as a general example of feedback systems. These observations may help in the analysis of other systems where the separate operations are unclear. Each of the parts of a control system can be identified: the feedback function is the spotter (plus his communication equipment); the output function is the gun (plus the powder, projectile, aiming devices and crew); the reference signal is the target (provided by higher command, the "Decision Making Entity"); and the comparator is the human (or a specialized device) that determines the size and direction of the error provided to the crew to adjust the aim of the gun. For a time scale fast enough to observe these events as they occur, analysis can emphasize any one of the components. A mathematical equivalent of each of their separate operations can be written. For a time scale so slow that the system has come to equilibrium, analysis concerns the operation of the entire system as a unit. This is equivalent to solving the equations for the controlled variable in terms of the reference signal (and system parameters). The result is the familiar form used to describe the operation of a closed loop feedback system.

PEOPLE

The preceding discussion has been in rather mechanical, abstract terms. Regarding the people operating the system, each one is primarily concerned with his own part in the detailed sequence of events, rather than the combined operations as a feedback system. Each person uses the skills needed for the immediate purposes. He selects and applies them as he understands of their function in the larger organization. He also coordinates them with his individual internal conditions and needs.

The commanding officer, using a time scale suitable to his needs, regards each combination of gun, crew and spotter as one of the parts of his output function. To him, each "rifle squad" (is this the correct term?) is a simple straight-through system: he assigns the target and the system performs. This can be considered as an S-R System with its performance improved by adding a negative feedback loop. This treatment, however, omits the events in between the "S" and the "R." For some purposes it is adequate.

TWO LEVELS BECOMES THREE LEVELS

The above is an example of a control system with two levels. Selecting suitable response times helps separate and identify the different levels.

By adding another level of command, we have a three level system. For the gun crew-spotter, the time scale would be of the order of minute, the time to fire a few shells. The "commanding officer," above, is concerned with the operation of his several guns. His time scale would be from minutes to hours, and, in turn, the higher commander works with larger scale tactics and strategy and even longer time scales. To him, the individual gun and crew with its assigned target is simply a tool to be used. He is concerned with larger scale results.

Consider, in passing, what happens when the chain of command is by-passed and higher order corrections are introduced too early!

MEMORY

Memory, expressed in several forms, is essential to the operation of his system. Some of the data are in the form of maps and instructions. Some are in the form of the aiming and firing mechanism of the gun. Some are in the form of remembered procedures and instructions. Some are in the form of remembered orders "from above." And so on.

In fact, the entire set of concepts, ideas, procedures and skills are all located within the memories of the participating individuals. Each must have available, as a minimum, those portions of the operation that apply to him. Perhaps this could be simulated with high speed computers and software, but the operating components must all be included in some form. Although the mechanical requirements are relatively modest, the memory capacity and programming to provide for automatic selection among many alternative actions is mind-boggling!

ATTENTION
Each participant must direct his attention to the assigned task, while "simultaneously" "paying attention" to several other variables, especially those in his immediate environment. This requires frequent shifting of attention among several perceived variables.

Regards Bob Clark

Date: Fri Feb 05, 1993 3:32 pm PST
Subject: intentionality?

[From Francisco Arocha] Rick Marken (930204.1430)

>If Baars starts talking like this about volition then I'm in real
>trouble (understanding-wise).

In some philosophical writings, ever since Brentano, the term "intentionality" has come to refer, in again some philosophical discourse (and maybe now in AI too), to the "aboutness" that Bill P. mentioned recently. However, it seems to me that Brentano (and some philosophers) confuse intentionality (a psychological category) with reference (a semantic category). So when people talk about intentionality they may be talking about reference. So the talk you want to listen to may not have anything to do with the psychological concept of intentionality, but with this "aboutness stuff". I don't know the rea

By the way, if you (a scientist) want to know about reference (not intentionality) you should check about the only philosopher (actually, he is a physicist) that scientists should read, that is, Mario Bunge. His Scientific Research and Treatise of Basic Philosophy are a must read. Of course, philosophers 1) never read anything he writes and 2) hate everything he has to say.

Francisco Arocha (not a philosopher).

Date: Fri Feb 05, 1993 3:41 pm PST
Subject: language

[Avery Andrews 930205.0800] (Oded Maler 930205)

>Just an impression: when you put on your professional linguist hat,
>you sound more like the motor-control professionals you so nicely
>quote and analyze. On the surface it looks like someone responding to

Could be so - someone else will have to do the job on me, I guess. The reason I think the Gross article missed its target is that (I recall) the major focus was on the idea that each lexical item had its own grammar, so that the transformational regularities that people kept claiming to find kept having embarrassing exceptions. The exceptions are still there, but when the transformations are defunct, or of greatly reduced scope, it's not clear that this means anything. (I would tend to follow Anna Wierzbicka in thinking that most of the apparent irregularities are semantically-based regularities).

As for `real language phenomena', well, generative ideas help with some of them, but not others. In Martin Taylor's terms, if I'm not butchering his stuff too badly, `syntactic' effects predominate in smaller units (words, phrases, sometimes whole sentences in special circumstances such as unviersity lectures), whereas with the larger units interactions start kicking in, & you need some different ideas to get anywhere.

I don't think anybody doing generative grammar thinks its all you need for a full theory of language - we just think that sharply defined and appropriate ideas about grammatical structure will be useful.

There are many schools in the subject - the predominant one, to which I do *not* belong, wants to come up with as `strong' an account of universal grammar as possible, so as to `explain' how children learn languages. I think this is overly hasty and presumptuous, and follow a minority crowd who want (a) more and better information about the grammatical phenomena of different languages (can circustantial modifiers modify & agree in case with semantically-cased marked NPs in Warlpiri - is it really true that in Bardi an NP without a case-marker can be split in two (interrupted by the verb), while one with a case-marker cannot be (b) better formalisms for basic grammatical description. E.g. the LFG formalism is better than transformational grammar in many respects, but it's got plenty of problems too, in that if you try to implement a basic grammatical description of a random language you will find yourself being forced to do all sorts of stupid and unpleasant things (tho far fewer than if you were trying to do a transformational grammar).

Summarizing, the things I regard as progress are of a fairly commonsensical nature - collection of basic information, formalisms that you can use a bit longer before get stuck (and write parsers for that run on low end PCs, to boot). But no sensible generativist claims to have a full theory of language - they'd follow Chomsky in saying that there just aren't the ideas around to make a full theory of language attainable, so its better to work on grammar, where some progress can be made. But I think it's quite likely that Chomsky is underestimating the chances for progress outside of grammar.

Avery.Andrews@anu.edu.au

Date: Fri Feb 05, 1993 4:28 pm PST
Subject: Motor programs; ongevity; flinging; conventions; etc.

[From Bill Powers (930204.1500)] Oded Maler (930205.1355 ET)--

I believe in motor programs, too, Oded. What I don't believe is that (in natural human control processes) they run open-loop, without constant checking that the program of desired effects, as perceived, is actually going according to plan. This checking doesn't have to be conscious, but I presume that the basic nature of a motor program is a planned series of perceptions, not of motor outputs. When a programmed plan is in process, of course there are patterns of motor output going on, which could easily be mistaken for a preset program of action. I will even go so far as to say that there is probably some sort of parametrically-adjustable oscillator or pattern generator involved in creating the output pattern of reference signals for lower-level systems.

I believe, however, that if one were to examine any such motor program carefully, preferably with instrumentation, one would see that there are always small deviations from the nominal program, and that there are continuous adjustments that prevent these deviations from becoming significant. If we actually examined the motor outputs themselves, the forces being applied to the body and environment, we would see that they deviate MARKEDLY from a regular pattern -- in exactly the way, and by the amount, and in the direction, that will maintain the pattern of PERCEIVED CONSEQUENCES nearly constant. If anyone actually did this experiment, I think it would be clear that the motor program does not govern outputs, but perceptions.

I should think that a simple version of such an experiment might be done rather easily. Can you think of an example to focus on?

Gary Cziko (930205.1434 GMT) --

My father smoked until his late 40s, drank rather freely at times, had a desk job in Chicago during the coal-burning days, and liked red meat. On the other hand, he stopped smoking, became a moderate drinker, and took to climbing mountains in his later years -- and spends hours sitting before a typewriter revising economic theory. His father, who lived to 96, smoked to the end, drank all he wished, ate whatever pleased him, and never took a lick of exercise on purpose.

Sorry. I don't think that there is much we can do to help us live as long as we would like to, or keep from living longer than we really want to. Or that there is any useful way of predicting how it will turn out.

Flinging suitcases into the air or failing to budge them: an interesting subject. It does seem that we prepare for anticipated control difficulties if they're expected to be outside the normal range (sometimes incorrectly). I think we can learn more about hierarchical control by studying such instances. Maybe the anecdotes will lead someone to do some actual experiments with this.

Bruce Nevin (930205.1041) --

>So please, you do likewise, and leave out reference to
>technical matters of linguistics while we work on the question
>you raised, namely what social agreements are, how we arrive at
>them, and how we learn them when they are faits accompli upon
>which those around us depend, and upon which we must learn to
>depend in like manner if we are to cooperate with those around us.

A most reasonable suggestion. What I'm most interested in is what a social convention (for anything) is that we can perceive it, and what it is we must perceive in order to know that we're experiencing one, or conforming to one.

For example, this must involve noticing patterns in what other people do and also experiencing the results of doing things differently. Clearly we don't imitate everything other people do; otherwise we'd all be doing exactly the same things. How do we choose which patterns to imitate and which to do our own way? What is it that we do when we see someone else deviating from a pattern we have accepted as the socially right one?

In another context, if we take communication to be a goal -- the transmission of meaning -- is there anything in general that we can say about the way two systems attempting to communicate will reorganize toward this end? What kinds of errors would develop when two people have different conceptions of the social agreement? Are there differences in the conceptions that could remain undetected even though communication appears satisfactory to both parties?

Oded Maler (930205.1740 ET) --

RE: intentionality:

>What is the semantics/denotation/aboutness of a reference
>signal (inside an individual)? It is the set of states-of-
>affairs in the world, such that when sensed by the individual
>will cause a zero error in the corresponding comparator.

To this we have to add there there is no unique actual state of the external world that will result in a perception that matches a reference signal. When such a match exists, the world can be in any state that still allows a specific function of the external variables to have the specified value. Basically the reference signal is "about" the perceptual signal, not the objective world.

And anyway, isn't this all sort of metaphorical? None of this aboutness or directness would happen if it weren't for the associated control system. If I put a probe on a reference signal and let you see the meter reading, you wouldn't have a clue as to what it was about, denoted, or meant.

Bob Clark (930205.1425) --

A very nice analysis with lots of interesting observations. One thing your examples about "synthetic" control systems shows is how crude control actually is when an organization tries to imitate individual control systems. But even an organization wouldn't think of computing how to aim the gun and firing it without looking to see where the shell landed.

One minor quibble:

>... the reference signal is the target (provided by higher
>command, the "Decision Making Entity"); and the comparator is
>the human (or a specialized device) that determines the size
>and direction of the error provided to the crew to adjust the
>aim of the gun.

When the commander says "Put a warning shot just in front of them," the aiming point is not the target, but a point that bears a specified relationship to the target position. So it's the relationship between the impact point and the target that is the reference signal, and it exists only in someone's head prior to the shot.

An added observation:

In order to adjust the gun position over repeated shots, the error must be turned into a new gun position. In order to get the final error as small as possible, you need a high loop gain. But if you have a high loop gain, an error of +50 yards would lead to a large correction, and the next error might be -500 yards. The solution is to use a slowing factor, such that only a constant fraction of the computed correction is actually applied on any one trial. In that way you can have high loop gain and accuracy, without instability of control. Same principle that applies in spinal control loops with transport lags.

Best to all, Bill P.

Date: Fri Feb 05, 1993 4:30 pm PST
Subject: Re: Intentionality

[From Rick Marken (930205.1500)] Oded Maler (930205 17:40 ET)

>should you attribute intentionality to a signal in a motor-neuron, or a
>thermostat?

It looks to me like there are two clearly different meanings of "intentionality" being used by people interested in understanding living systems:

1. Aboutness -- call this "intentionality - a"
2. Purposiveness -- call this "intentionality - p"

I would only use "intentionality-p" to describe the behavior of the entire negative feedback control loop; perceptual signal, error (motor) signal, output variable, input variable. So I would attribute intentionality to a properly functioning thermostat -- because it really is intentional (in the rigorous sense of PCT); it is controlling its perceptual signal (the voltage across the themocouple) relative to a reference voltage (set by you when you set the "temperature" of the thermostat).

I would use "intentionality-a" to describe variations in the perceptual and reference signals in the control loop; variations in the perceptual signal are "about" the variable or the function of several variables "out there" that result in variations of the perceptual variable; the same is implicitly true of the reference signal. The voltage across the thermocouple in the thermostat is "about" temperature.

PCT has interesting things to say about both intentionality-p (which is probably the main concern of your basic PCTer) and intentionality-a. PCT says that intentionality-p is the control of perception. This means that you cannot know what an intentional system is "doing" unless you know what it is trying to perceive. Knowing what an intentional system is trying to perceive (from the point of view of an observer) seems like a matter of trying to figure out what the system's perceptions are about (intentionality-a). But neither the system nor the observer can see beyond his or her perceptions; so there is no way for the observer (or the system) to know what their perceptions are "really" about. So for PCT, determining the intentionality-a of the system means learning to perceive what the system is perceiving (or learning which aspect of the observer's own perception is a correlate of the system's perception).

Best Rick

Date: Fri Feb 05, 1993 4:44 pm PST
Subject: motor programs

[Avery Andrews 930205.1100] (Bill Powers (930204.1500))

> but I presume that the basic nature of a motor program is a planned
> series of perceptions, not of motor outputs.

This is a crucial thing to emphasize (when talking to people like Randy Beer, for example), since it's a simple idea that the motor control community really doesn't seem to have gotten a grip on. The other thing to emphasize is that a `motor program' doesn't, and mostly surely usually isn't, represented as anything vaguely resembling an explicit symbolic plan (except perhaps in certain kinds of extreme cases, such as when one gets stuck in rock-climbing).

What there is instead is nonlinear circuit elements that spit out the right wave-forms, with more or less dependence on what comes back from the environment (the continuous, dynamic version of old-fashioned response-chaining, I guess). This is where I suspect the Kugler & Turvey stuff may be quite useful, in spite of the silly things they say about control theory.

Avery.Andrews@anu.edu.au

Date: Fri Feb 05, 1993 7:47 pm PST
Subject: Iterative vs. Continuous Control

[from Gary Cziko 930206.0225 GMT]

Bob Clark (930205) on BALLISTICS says:

>TIME SCALE & "CLOSED LOOP" VERSUS "OPEN LOOP"
>Analysis is influenced by the time scale selected. When times of the
>order of seconds are of interest (approximately the time needed for the
>projectile to arrive), there is no control without self-guidance. Here
>open loop analysis applies. Events are followed around the loop without
>treating the system as a whole. When events are examined in terms of the
>time for firing the gun several times (several minutes), closed loop
>analysis applies to each firing of the gun as the assigned target is
>followed. Assuming the necessary components are present, either closed or
>open loop analysis may be suitable according to the time scale of interest.

This ballistic artillery example is very interesting. You use it to argue that a open-loop or closed-loop analysis depends on the time scale. But in the artillery example, no matter how quick the feedback from the spotters, it still seems quite different from the continuous control systems that PCT has introduced me to. In a "real" continuous control system, perception, comparing, and acting are all taking place AT THE SAME TIME. And with the proper slower, a very high loop gain can be used for keep error low while maintaining stability. In your ballistic artillery example, we have instead an interative control process which seems more like S-R chaining with a reference level, like the TOTE system described by Miller et al. And I don't see how you can have a high loop with such a system. If the the spotters say that the shell missed the target by falling 100 meters too far, the gunner cannot use a loop gain of -10 and aim 1,000 meters shorter the next time around.

So isn't there really a qualitative difference between the iterative control you describe and the continous control of engineered and living controls systems? And doesn't this involve more than just the time scale?--Gary

Date: Sat Feb 06, 1993 11:38 am PST
Subject: Ballistics & PCT, progress

[From Dag Forssell (930205 11.30)]
Gary Cziko 930206.0225 GMT BALLISTICS & PCT - RKC

>In a "real" continuous control system, perception,
>comparing, and acting are all taking place AT THE SAME TIME.

Nothing happens AT THE SAME TIME in our loops. Your radio contains some electronic servo amplifiers where the signals travel around the loop with something approaching the speed of light. (I am somewhat ignorant here. Do not believe that electronic signals in wires necessarily travel with quite the speed of light. Happy to be enlightened). As a result, your amplifier is able to follow and amplify a reference signal that varies with radio frequencies. I forget the number, but it has been stated on this net (discussing beer's cockroaches)? by what speed neural signals travel. ??? feet per second. As a result, the minimum reaction time in a neural loop of five centimeters total length is ?? milliseconds. The minimum reaction time in a human loop of four feet (from your calf to spinal cord)?? would be .?? seconds, enough to counter a stumble, but certainly not instant.

PCT demo: (not even a rubber band)!

Two persons stand facing each other with index fingers opposite each other. Person one represents the (weak)! audio signal as received by (Gary's ham) radio. Person two is the amplifier who will faithfully reproduce the signal at a power level millions of times stronger. Person two therefore imagines that her hand is in a vat of molasses (offering resistance, requiring effort).

Gary (person one) has now tuned in the station and begins to walk slowly sideways while moving finger up and down in unpredictable (but rather slow) pattern. Dag (person two) follows Gary's finger quite well with his finger while grimacing and grunting to indicate the tremendous effort expended to overcome the resistance of the circuitry that follows and still send out a powerful signal.

The only difference between the two people and your radio is the signal traveling speed and of course the kinematics etc.

>So isn't there really a qualitative difference between the

>iterative control you describe and the continuous control of
>engineered and living controls systems?

No. Only a quantitative difference. As long as the components of the system are dedicated to their task and function in a dependable way, they are still control systems. (This proviso is necessary to distinguish between control systems and social (control) systems. This proviso is unrealistic with people in a community or organization, which is why we say that there are *no* social control systems).

Progress report: Thursday, Christine and I gave a talk:

Perceptual Control Theory
a complement to the
Deming Management Philosophy

to the Los Angeles Deming Users Group. It was well received. We shall now edit a two hour videotape of this presentation.

Best, Dag

P.S. Glad to see that you added the new references to CSGintro, Gary.

Date: Sat Feb 06, 1993 12:37 pm PST
Subject: AT THE SAME TIME

[From Dag Forssell (930205 12.30)]

Better than "at the same time" is CONTINUOUS.

The reason for the failure of many to understand control is - I believe - an attempt to think of the control process as a set of discrete, complete happenings.

1 set a goal
2 execute a program designed to achieve the goal
3 perceive the result
4 compare and determine your level of satisfaction.

The electronic amplifier demo makes it clear that this is not happening in that case. Consider the alternative:

1 Memorize a seconds worth of wiggly lines
2 write a seconds worth of wiggly lines
3 perceive the result
4 compare and determine the harmonic distortion

An excellent demonstration of CONTINUOUS is the Marken spreadsheet model. When you reset one of the top logical relationships, you make a major step change in the reference signal.

The control hierarchy does NOT execute a program designed to achieve any goal. Error signals arise.

Pay close attention to the iterations the spreadsheet program goes through. Do Not set it for 100 calculations all at once, which gives you the impression that the hierarchy quickly and effortlessly adjusts.

Do set it for single iterations or perhaps 10 at a time. Observe the small changes that take place in response to the error signals. Notice particularly that the goal is far from reached and the error signal is still large after a few iterations.

The spreadsheet model shows you how a control system incorporates references, disturbances and perceptions "at the same time" and continuously, but not necessarily in zero time. The competing models cannot handle all the things that happen "at the same time" so they think in terms of chains of cause-effect.

I was fully trained in the Swedish Army Artillery. A behaviorist might listen to Bob Clark and hear him say that this is a chain of cause-effect happenings. We in PCT notice the *multiple* iterations required to arrive at the target and can see the similarity with the iterative calculations of the Marken spreadsheet. We can see that the difference is quantitative, not qualitative, since we see the error signals at work, pulling in some (hopefully correct) direction and know that the process works well even without perfectly planned and executed output functions.

PCT demo:

The demo I just described could be done through a dirty or misty window, which would leave a trace. You could make person 2 write something in chinese if you wanted. Obviously person 2 could work against friction. If you limit the field of vision, person 2 would not know ahead of time (of course) and not afterwards either what she wrote. (Backwards, upside down).

Again, Dag

Date: Sat Feb 06, 1993 12:45 pm PST
Subject: Re: demo of linguistic nastiness, error control

[Martin Taylor 930206 15:40] (Rick Marken 930202.2100)

> error can be controlled; but when it is, it
>must be represented as a perception (of the error signal) in a
>control system that has its own error signal that indicates the
>discrepency between what it intends and what is the actual level of
>perceived error.
>
>Right?

Sounds like a recipe for an exotic kind of conflict. If ECS A has an error, it will produce an output until that error is zero. If ECS B uses the error of A as a perceptual signal that it tries to control to a non-zero value, either it fails because A succeeds, or it succeeds because A fails to keep its error at zero, or both fail. Isn't this the normal result of conflict?

Sorry if this is superseded by other messages in the last 3 days. I am trying to call in from home, and the lines are normally not good enough to allow me to try reading and aswering mail. There are 110 messages waiting, so I'm trying to answer while the line seems to work, without reading ahead to see if the point has already been made.

Martin

Date: Sat Feb 06, 1993 4:26 pm PST
Subject: LEADER/FOLLOWER - RKC

>Bill Powers (930202.0138)

Your more complete description the Leader/Follower demo is helpful. And I would join in the claim:

>that the minimum possible time required for this swap is longer than the
>time taken to change any lower-level control process.

I would like to point out that the "lower-level control process" consists of tracking the Leader's finger. This requires control of muscle variables, position variables, and time variables.

The Follower uses some of his Fifth Order Skills. From Bob Clark (921205). He has formed recordings from observations and, perhaps, tracking experiences. He (DME) can select one that may produce acceptable results. He (DME) uses this to provide reference levels to produce his tracking movements. As the pattern changes, different recordings are needed. It takes more time for these changes than it takes for lower level (muscles, positions) changes. To the Follower, this is still the Tracking Demo.

The Leader also has a supply of recordings available from his experiences, etc. His assignment as Leader calls for him (his DME) to select one to to be tracked by the Follower.

Bill, you report the results when

>one of the people simply changed roles without warning the other and
>without any external signals.

I find it interesting that you speak of "roles" in this connection. Where are your higher level systems? To me, "role" is a Sixth Order concept involving a combination of Sixth Order Perceptual Variables. Thus terms like "Leader," "Follower," and "Role" belong to Sixth Mode of Sixth Order.

Warning! If your subjects are unfamiliar with participating in such demonstrations, there can be some unexpected side effects. For example, some people avoid the role of Leader. Being a Follower may be acceptable, but being the Leader introduces some intrinsically different perceptions. The particular behavior depends, of course, on the specific individual. As switching becomes faster, the participants may become confused and conflicts (internal) at Sixth Order may develop.

I suggest you examine your own -- remembered -- internal experiences when you have been a participant in this demo.

Bill, still from Bill Powers (930202.0138), I am surprised by your reaction to my, Bob Clark (930131), remark. You remark:

>you seem to be looking for levels that will apply to "psychological"
>aspects of a person, to explain the how and why of that person's
>behavior."

This suggests that I BEGIN by selecting "psychological" aspects, FOLLOWED BY searching for lower order systems (variables?) that might fit. To the contrary, I begin with the lower order variables, defined by Skills and lower Order variables. Thus I look for perceptual variables that use combinations of selected Skills (including their related lower order variables. With a rather large assortment of these perceptual variables, the question is one of assigning useful labels. Labels are needed to facilitate their selection and application, both for use as sources of sets of reference signals and for communication. Labels are preferred that will be generally understood and thus communicate to more people.

Bill, I am basing my analysis on your very important observation that BEHAVIOR IS THE CONTROL OF PERCEPTION, and that perceivable variables are the heart of the structure. I may well have overlooked some important aspects of the situation -- I am sure you will point out where my suggestions can be improved.

I plan to work up a more complete discussion of this as well as comments on your parallel description of your approach to "definitions of levels."

Regards, Bob Clark

Date: Sat Feb 06, 1993 4:48 pm PST
Subject: sequential vs lagged control; control of error

[From Bill Powers (930206.1400 MST)] Dag Forssell (930205.1130) --

There's a subtle difference between "sequential" and "lagged" control. Bob Clark gave an example of truly sequential control: lob a shell, wait for the spotter to see where it lands, wait for the spotter to send the message back to the gun site, lob another shell, etc.

Lagged control is like aiming a fire hose. The water shoots through the air and lands somewhere. The fireman is watching where the water lands, and corrects his aim according to the error between perceived and intended landing spot. There is water continuously flowing and continuously landing, and the fireman is continuously monitoring the landing spot. There is always water leaving the nozzle at the same time that water is landing on the fire, at the same time that the fireman's eyes are seeing the water landing, at the same time that the fireman's muscles are altering the aim of the nozzle. The processes in various parts of the loop are all going on at the same time, literally simultaneously -- even if it takes two or three seconds for any one drop of water to fly through the air and land on the fire, and a hundred nanoseconds for the image of the water landing on the fire to reach the fireman's eyes, and 50 to 200 milliseconds for the image to be converted into a perceptual signal, and an error signal, and a new muscle tension.

The second case is the most common in human behavior, although there are valid examples of the first (corresponding by email, for example).

Many analysts of human behavior have confused sequential control with lagged control. They assume that while a stimulus is occurring, everything else in the control loop is on hold until the stimulus finishes its pattern. Then, with the stimulus input finished, the response commences, goes through its pattern, and stops. At that point the effect of the response alters the stimulus conditions, with neither stimulus nor response occurring. Finally, the next stimulus occurs and the sequence begins again.

Even inside the nervous system this same erroneous image seems to be used. A neuron fires, sending an impulse along a fiber to its end, where the impulse triggers off the next impulse in line. The maximum number of input-output events per second therefore seems to depend on the time it takes for an impulse to travel through the nervous system to a muscle.

In reality, there can be anywhere up to 10 or so impulses travelling along the same nerve fiber at the same time (length of path, say 0.5 m., divided by speed of travel, say 50 meters per second, times impulse frequency, 1000 per second or more). The maximum number of input or output pulses per second is set by the maximum impulse rate, regardless of transit time through the nervous system. If you count redundant paths carrying similar information, the maximum rates are even higher than that. For a complete analysis, ask Martin Taylor.

This confusion is the result of trying to describe a closed-loop process in words. Using words we can say only one thing at a time. We can't be talking about input processes while we're also talking about output processes and the processes in between, or the effects going on in the external part of the loop. So language forces us to describe first the input, then the comparison, then the output, then the effect on the environment, then the effect on the input again, as if this were a sequence of mutually-exclusive events. If one lets words dictate thought, the mental image of the process will have the same sequential nature, leading to incorrect analyses and failed predictions.

Martin Taylor (930206.1540) --

Once again we gang up on poor Rick Marken (but more so on Mark Olsen). I agree with you: controlling error doesn't mean anything. The error signal is part of a control process. To control something means to act on it to bring it to an arbitrary and by implication adjustable state. Control systems don't create an arbitrary adjustable amount of error signal: they create only zero error, as nearly as possible.

I think that Rick speaks confusingly because he's really talking about error in terms of two different logical types of variables. It isn't that he doesn't know what he's talking about.

Suppose you want to build an adaptive control system. You start with a plain control system with, for our purposes, one adjustable parameter, output sensitivity. With a given value of the parameter the control system will be able to control with some degree of accuracy, stability, and speed, given some standard external feedback link. When the parameter has a low value, the error signal will fluctuate a lot as random disturbances come and go. Gradually raising the value of the parameter, you will observe that the average amount of error fluctuation decreases up to some value of the parameter, and then begins to increase again as the control system approaches instability (too much loop gain for the lags around the loop).

An "optimizing control system" is something of an oxymoron; either the control system controls, or it optimizes. To get optimization of the simple control system above, we must add a SECOND control system with its own perceptual signal, reference signal, comparator, error signal, and output signal. This is where the confusion about "controlling error" as Rick means it begins.

From the standpoint of the optimizing system (which is really just an ordinary non-adaptive control system), the controlled variable is some aspect of the error signal in the first system. This aspect can't be the same one that's significant in the first control system, namely amplitude, because as Martin comments, that would simply lead to "an exotic kind of conflict." So the second control system must perceive a different aspect of the error signal that is not part of the first control system's operation.

An appropriate aspect would be _average squared amplitude_ or _average absolute amplitude_. This is what we mean when we say the "error signal is too large." We mean it is too large for both positive and negative excursions, over some period of time. We can't mess with the instantaneous amplitude, because that's what's making the first control system work.

So the second control system, the one doing the optimizing, monitors the average squared or absolute error signal in the first control system. It compares this with a reference signal, which might as well be zero, and converts the error signal into an output signal on which the output sensitivity parameter of the first control system depends.

We have a problem here, because an increase in the parameter might either decrease or increase the average squared error; for small values of parameter, increasing the parameter will decrease the average squared error, but past a certain (unpredictable) point, increasing the parameter will increase the average squared error because the first system starts to become unstable. If the feedback around the second system is negative for small values of the parameter, it will become positive when the parameter exceeds the critical value. Then increasing the parameter more will cause the average squared error to get larger, not smaller.

If we make the second control system into a reorganizing-type system, this problem of the feedback switching sign will not be important. The output of this system now simply increases slowly with time (or decreases, depending on the starting conditions). If the second system's error signal is decreasing over time, the change in the output signal (and the first system's output parameter) simply continues. If, however, the second system's error signal begins to increase with time, then the output signal switches from slowly increasing to slowly decreasing or vice versa, depending on which it was before. We can also make the rate of increase or decrease depend on the absolute magnitude of the second system's error signal, so the changes will get slower and slower as the second system's error signal becomes smaller.

Now the second system will raise the parameter of the lower system as long as the average squared error signal in the first system continues to get smaller. If that error signal starts to get bigger, the second system will switch to decreasing the parameter. In the steady state, the parameter will be slowly increasing and decreasing, always overshooting the right value but immediately reversing its course to head back toward the right value.

This second system can be left "on" all the time. If the environment changes its properties, the second system will simply re-optimize the first system automatically.

The combination of the first control system and the second one now constitutes an "adaptive control system" made of two systems neither of which can adapt itself. The right way to speak of this is to say that the second control system adapts the first one for optimum performance. To an outside observer who doesn't know the details, it may seem that there is a single "self-adapting" system here. But it is really two independent control systems, even if the two types of control are cleverly combined into a single system. And neither system adapts _itself_.

The most important point is that while the second system is indeed monitoring the error signal in the first system, what it is monitoring is not the same aspect of that error signal that is important in the operation of the first system. That's the only way we can have "control of error" without creating conflict, or without achieving nothing more than a change in the reference signal of the adapted system.

Best to all, Bill P.

P.S. I'll be off the net for the next 10 days unless I come
across an amenable computer somewhere (I'll be seeing Rick and
Dag in a week, so who knows?).

Date: Sun Feb 07, 1993 8:03 pm PST
Subject: more on rubber band demo

from Ed Ford (930207:1055)

Bruce - The material you asked for has been mailed...

Gary - the books have been mailed.

To all paid-up CSG members - Closed Loop will be mailed Tuesday, Feb. 9th.

To all - More on the rubber band demo. I spent the last week in Durand, Michigan (between Flint and Lansing) training 32 teachers, counselors, administrators in control theory (among other things). In showing the rubber band demo where a teacher held two knotted rubber bands stretched out, with the knot directly in front of her chest. She had to keep the knot right at the tip of my pointed finger that was moving. When I asked the participant to watch the action of her right hand instead of the knot, I began watching her eyes. I could see her eyes occasionally sneaking a look at the knot. Thus, she was able, by sneaking an occasional look at the knot in relation to the dot, to achieve her goal, but with less efficiency. I wanted to force her to just watch her actions, so I got someone else to take my place at moving his pointed finger in front of the knot. Then I took a cardboard about 12 x 8 inches, and placed it between her eyes and in front of the knot. Now she couldn't see the knot and the tip of my finger and her ability to little glances was eliminated. Her ability to keep the knot over the dot became far more pronounced, in fact she couldn't do it at all. That demonstrated clearly that we need feedback to achieve a goal and that watching our behavior has nothing to do with controlling a variable.

Another thing I found is when you use an actual example to explain how control theory works, the concept becomes a lot easier to understand. For example, a little boy in class talks to a friend while the teacher is speaking. The boy perceives talking to a friend more important that listening to what the teacher is saying. He makes a choice to talk and asks a boy next to him about another friend in school. The child reacts by talking back to him. He is satisfying his goal. The teacher has several goals, one to teach and another for the class to be quiet and attentive. She perceives the little boys talking which is acting as a disturbance to what she wants, namely a quiet class. The teacher then acts as a disturbance to the boys by asking them to stop talking. And so forth.

This kind of example really helps teachers not only understand PCT but also they could more easily see its relevance to the various techniques we were teaching.

Best, Ed

Date: Sun Feb 07, 1993 8:07 pm PST
Subject: Re: Motor programs; ongevity; flinging; conventions; etc.

[Martin Taylor 930207 13:20] (Bill Powers 930204.1500)

Sorry you'll be away when this arrives, but anyway:

>Flinging suitcases into the air or failing to budge them: an interesting
>subject. It does seem that we prepare for anticipated control difficulties
>if they're expected to be outside the normal range (sometimes incorrectly).
>I think we can learn more about hierarchical control by studying such instances.
>Maybe the anecdotes will lead someone to do some actual experiments with this.

No experiment, but another anecdote of a different kind, leading me to think the problem has nothing to do with the expectation being outside the normal range. I told this one to Rick privately, but it might be worth making public.

I wash dishes using baking soda rather than detergent (works better, and is said to be kinder to the environment). The method is to wet the dish and a finger or two, dip the finger into a pot of baking soda, and rub it over the dish, then rinse. For a long time we have kept the baking soda pot at the end of the windowsill behind the sink, but about ten days ago, my wife moved it to the middle of the windowsill (and of the sink). Even though I know this, and can see the pot in the middle of my field of view, the first move of my hand is always toward where the pot used to be, and the movement nearly reaches the windowsill before I correct it (with appropriate mental swear words about my continued stupidity in not remembering the pot had moved).

We usually talk about ongoing control, and we have talked vaguely about controlling not so much for perceptions based strictly on sensory data, but for perceptions based on the difference between sensory data and expected sensory data. I think what happens here, and with flinging suitcases, is that a reference signal is moved abruptly, or a perception that was not being controlled is brought under real-world control, and the initial output is derived from an error signal based on imagined perception. Usually that output is appropriate, more or less, leaving only a little correction for ongoing control.

Martin

Date: Sun Feb 07, 1993 8:44 pm PST
Subject: anticipation

[Avery.Andrews 930208.0914] (Martin Taylor 930207 13:20)

It would be excellent if this sort of `anticipation' (control of imagined perceptions) could be built into a vivid computer demo, since the limitations of feedback control of errors that have already occurred are a familiar theme in the anti-feedback literature.

Avery.Andrews@anu.edu.au

Date: Sun Feb 07, 1993 8:46 pm PST

[From Rick Marken (930207.1600)] Martin Taylor (930206.1540) --

Boy. You guys just won't let me get away with ANYTHING. You are right; if the error signal itself were controlled at anything other than zero (which the control loop keeps it at anyway) then there would be terrible conflict.

Bill Powers (930206.1400 MST) agrees with you and says:

>Once again we gang up on poor Rick Marken

Just wait 'till he gets here. Har har!

Nevertheless, I must admit that Bill's comments on this issue (control of error) were very helpful and might help me get the spreadsheet reorganizing system working.

Martin Taylor (930207 13:20) --

>Even though
>I know this, and can see the pot in the middle of my field of view, the
>first move of my hand is always toward where the pot used to be, and the
>movement nearly reaches the windowsill before I correct it

and Avery.Andrews (930208.0914) suggests

>It would be excellent if this sort of `anticipation' (control of
>imagined perceptions) could be built into a vivid computer demo,

A demo of something that seems somewhat similar to what Martin described (but not necessarily the same thing) is one of the Powers, Clark & MacFarland "portable" demos -- and it could easily be turned into a computer demo. Just have the subject track your finger with her finger (I just did this with my daughter) as it makes a regular pattern (an approximately 8 in diameter circle seems to work nicely). Move your finger at the rate of about 1 cycle per second -- slow enough for good control but fast enough so that knowing the circular movement pattern really helps. Then stop your finger at an unpredictable time. Your subject's finger not only takes a while to stop (about 1/2 sec) but while it is moving it is tracing out an obvious CURVE, even though there is no target present to track. So the movement after the signal to stop is still controlled relative to a reference circular movement. There is "anticipation" that the target finger will not only continue to move but that it will continue to move in a circle. (I put "anticipation" in quotes because this could be modelled with any explicit computation of predicted target position at all; the model just controlling a higher order variable that might be called "relative circular motion").

Now do the same thing but use irregular mvoements of your "target" finger. Try to move you finger at about the same rate at which you were moving it to make the circle. I did it by writing out some words in the air. Now, when you stop the finger you will find that the subject moves very little after the stop. This is because (in theory) the tracking is now being done at a lower level; If target movements are sufficiently unpredictable, there is nothing the subject can control except the distance between target and finger (a configuration). So there is no change in variable to be controlled when the target finger stops; the distance between target and finger is all that must still be controlled. But when the target was a circle, the stopped target changes the variable controlled from "circular pattern" (probably an event level perception) to no pattern.

Anyway, it's a nice way to spend a few minutes with your kids. My daughter got a kick out of seeing her finger keep moving in a curve after mine stopped; even though she was trying very hard NOT to let that happen. I didn't mind humiliating her in this way because she keeps beating me at every computer game I've got.

Best Rick

Date: Mon Feb 08, 1993 1:25 am PST
Subject: Motor programs, intentionality and time-scales

[From Oded Maler 930206.1000] Bill Powers (930204.1500):

* I should think that a simple version of such an experiment might
* be done rather easily. Can you think of an example to focus on?

I don't know if that is what you meant, but I think it's about time that you teach your Little Man to draw and write. A motor program that draws "A"s or triangles will be a very impressive demonstration of "your" notion of a motor-program, especially if the size, the location (relative to the shoulder) can also be adjusted by higher references, and there is a lot of noise between the levels. Maybe in order to really impress the establishment you will need more degrees of freedom, but I think the current system is enough for a start. I think there were experiments concerning how people write when you attach springs and weights to their limbs.

Have fun on the mountain, meanwhile we ask Rick to build us the golden veal (a stupid biblical joke, probably with the wrong choice of words, but I could not resist).

* ---
* Oded Maler (930205.1740 ET) --
*
* RE: intentionality:
*
* >What is the semantics/denotation/aboutness of a reference
* >signal (inside an individual)? It is the set of states-of-
* >affairs in the world, such that when sensed by the individual
* >will cause a zero error in the corresponding comparator.
*
* To this we have to add there there is no unique actual state of
* the external world that will result in a perception that matches
* a reference signal. When such a match exists, the world can be in
* any state that still allows a specific function of the external
* variables to have the specified value. Basically the reference
* signal is "about" the perceptual signal, not the objective world.

So what is the perceptual signal about?

* And anyway, isn't this all sort of metaphorical? None of this
* aboutness or directness would happen if it weren't for the
* associated control system. If I put a probe on a reference signal
* and let you see the meter reading, you wouldn't have a clue as to
* what it was about, denoted, or meant.

But this metaphor is the convention undelrying the the use of language as a medium for communicating meaning. If my thoughts, which are presumably, neural patterns, are not "about" anything external, how can I communicate with you? Should I tell you "You know, neuron X254 is oscillating in this waveform" and you'll answer "I know that feeling, but you are wrong, it should oscillate completely differently".

[Rick Marken (930205.1500)]

It looks to me like there are two clearly different
meanings of "intentionality" being used by people
interested in understanding living systems:

1. Aboutness -- call this "intentionality - a"
2. Purposiveness -- call this "intentionality - p"

I would only use "intentionality-p" to describe the
behavior of the entire negative feedback control loop;
perceptual signal, error (motor) signal, output variable,
input variable. So I would attribute intentionality to
a properly functioning thermostat -- because it really
is intentional (in the rigorous sense of PCT); it is
controlling its perceptual signal (the voltage across the

themocouple) relative to a reference voltage (set by you when you set the "temperature" of the thermostat).

I would use "intentionality-a" to describe variations in the perceptual and reference signals in the control loop; variations in the perceptual signal are "about" the variable or the function of several variables "out there" that result in variations of the perceptual variable; the same is implicitly true of the reference signal. The voltage across the thermocouple in the thermostat is "about" temperature.

PCT has interesting things to say about both intentionality-p (which is probably the main concern of your basic PCTer) and intentionality-a. PCT says that intentionality-p is the control of perception. This means that you cannot know what an intentional system is "doing" unless you know what it is trying to perceive. Knowing what an intentional system is trying to perceive (from the point of view of an observer) seems like a matter of trying to figure out what the system's perceptions are about (intentionality-a). But neither the system nor the observer can see beyond his or her perceptions; so there is no way for the observer (or the system) to know what their perceptions are "really" about. So for PCT, determining the intentionality-a of the system means learning to perceive what the system is perceiving (or learning which aspect of the observer's own perception is a correlate of the system's perception).

Int-p, I agree, is not problematic when you stay inside the organism. It resides in the reference signal. Int-a, "aboutness" is very important as I indicated above unless you adopt solipsism, deny the fact that your thoughts and perceptions are about something, and abandon the illusion of language as meaning carrier (rather than social ritual). I thought a bit about the quantitative differences between the aboutness of, say a signal in a thermostat, a specific sensory neuron in my body, and the complex cell assembly in my brain that is presumably "about" an abstract entity such as, say, 'marken@aero.org'. It is true, as Bill noted above that the set of state-of-affairs that correlate with an internal signal is infinite, yet it is not the whole set of all possible situations in all possible worlds. The signal in my home theromstat is about the temperature in my house, and its value in this signal, represents some fraction of the set of all possible worlds. Of course, its aboutness could be changed, by damaging the circuit, by moveing the thermostat to your house such that its denotational value will change.

Sorry, I have to go, now, I'll complete this line of thoughts next week.

Just two short ones:

[Gary Cziko 930206.0225 GMT]

So isn't there really a qualitative difference between the iterative control you describe and the continous control of engineered and living controls systems? And doesn't this involve more than just the time scale?--Gary

It is "just" a time-scale, but it looks to us qualitative when this time-scale crosses the time-scale on which our concious thining works.

[Avery Andrews 930205.0800]

>Just an impression: when you put on your professional linguist hat,
>you sound more like the motor-control professionals you so nicely
>quote and analyze. On the surface it looks like someone responding to

Could be so - someone else will have to do the job on me, I guess.

Just to check whether you can really qualify as a prototypical devil.

--Oded

Date: Mon Feb 08, 1993 8:03 am PST
Subject: Clinical Strategies

[FROM: Dennis Delprato (930208)] Ed Ford (930207:1055)

>Another thing I found is when you use an actual example to explain
>how control theory works, the concept becomes a lot easier to
>understand. For example, a little boy in class talks to a friend
>while the teacher is speaking. The boy perceives talking to a
>friend more important that listening to what the teacher is saying.
>He makes a choice to talk and asks a boy next to him about another
>friend in school. The child reacts by talking back to him. He is
>satisfying his goal. The teacher has several goals, one to teach
>and another for the class to be quiet and attentive. She perceives
>the little boys talking which is acting as a disturbance to what
>she wants, namely a quiet class. The teacher then acts as a
>disturbance to the boys by asking them to stop talking. And so
>forth.
>
>This kind of example really helps teachers not only understand PCT
>but also they could more easily see its relevance to the various
>techniques we were teaching.

Ed, in the above not-so-hypothetical case, are you implying that one solution that follows from PCT is the teacher saying something, such as, "Billy and Harold, please refrain from talking"?

I am interested in identifying concrete clinical implications of PCT--what might be called principles, strategies, and procedures.

Dennis Delprato psy_delprato@emunix.emich.edu

Date: Mon Feb 08, 1993 9:01 am PST
Subject: Re: control article

[Hans Blom, 930208] (Rick Marken (930204.0900))

> What do they [Levine and Loeb] mean by 'new types of
> control'? There is only one type that I'm aware of --
> maintenance of a perceptual variable at a specified
> reference level in the context of variable disturbances.

All the time? Sometimes just the end point seems to matter (gun control), sometimes just one intermediate point, such as in high jumping. In gun control you might choose between a flat shot or a high shot with equal results. Bill Powers (930204.0900):

> Human beings hardly ever control the "full trajectory."

If that is the case, 'new types of control', which do not try to maintain minimum error between a reference value and perceptions at all times, might provide superior performance in some cases. Or greater ease. When I fly to New York, I (attempt to) control my destination, but in the plane I have to trust the pilot. Part of my trajectory will be, as far as I am concerned, ballistic.

> Levine and Loeb [...] seem to have forgotten to mention step one in the
> study of biological control systems -- the test for the controlled variable.

The controlled variable? What makes control in organisms so difficult to study is the simultaneity of a great many different ongoing goals, whose importance may, moreover, fluctuate from moment to moment due to influences beyond our control and usually beyond our knowledge. Only in the simplest of experiments one variable can be considered to be controlled, if at all. 'Keep your finger pointed at the knot', you ask. But the subject also has to control the upright position of her body and otherwise keep all sensory channels open, if only to hear you say 'you can stop now'.

> How can they compare the performance of a living system to
> know[n] "control schemes" if they don't know what ... the
> living system is controlling? Do they explain how they know
> what variables are controlled by the high jumper that they
> mention? If so, how do they know?

To me, this seems to be a clearcut case: a high jumper wants to jump as high as possible, period. An objective measure is provided to test that performance. All else is unimportant (within limits, see below). What more can you ask for? There is no prescribed tra- jectory to be followed; a new world record often is an unpreceden- ted experience for the jumper.

(Bill Powers (930204.0900))

> Human control systems are pretty close to the design limits
> set by the materials used. It's possible, for example, for
> an arm muscle to pull itself loose from its attachments to
> the bones, if feedback is lost and an energetic movement is
> attempted. Even with an intact set of control systems,
> tendons and muscles can be ripped loose if an emergency
> situation results in sending abnormally large reference
> signals to the spinal motor neurons.

Human control systems normally function well within design limits. We have very little experience with operation near those limits: pain effectively causes us to stay away from them. But pain is carried by slow nerve fibers; in emergencies the experience of pain may arrive too late to prevent harm. Is a case where 'tendons and muscles can be ripped loose' really an indication of 'an intact set of control systems'? I consider that to be pathology, a control system gone haywire, operating beyond its design limits. I would maintain that one of the most important of an organism's objectives is, at all times, not to seriously damage itself. But that cannot be formalized by control in the usual sense of the word, that is a perception following a reference signal. The control system is operating under constraints, i.e. it must _stay away from_ certain experiences with a high probability of success (catastrophy theory studies such problems). Short term goals are rarely important enough to jeopardize long term goals, which need an intact organism.

> Even the muscles work differently from the servo motors
> that engineers use. They don't apply forces directly, but
> by shortening the contractile elements in the muscle to
> alter the resting length of the series spring component.

That is also what Levine and Loeb maintain, and they show how difficult it is to reach top performance with such 'difficult' actuators.

> The reason a human being can't perform a mathematically
> optimal jump is simply the rocket problem: you would need
> to produce an impulse of muscle force of zero duration and
> infinite amplitude. That would hardly be a feasible
> solution for a servomechanism, either.

Impulses are not required, step functions will do nicely. After all, a trainer just want to study the peak performance that a real individual is capable of given her motor equipment, and search for whatever means there are to teach her to fire her nerves in such a way that this peak performance is reached.

> The "feedback too slow" argument turns up even here,
> doesn't it? Actually the speed of feedback in human control
> system is just right -- to explain the behavior we see.

Levine and Loeb do not say that feedback is too slow; bang-bang control requires very accurate timing. They say that when the need for performance becomes extreme, protection mechanisms are requir- ed to prevent muscles and tendons from being torn loose. Feedback from those protective sensors would probably be too slow if train- ing did not slowly familiarize the high jumper with the sensations that they provide (1). This is much like walking as closely to the abyss as you dare without risking the damage that a fall would cause (2). The fall would provide you with feedback, of course, but you wouldn't want _that_ feedback, would you?

1) Much of psychotherapy seems to serve the same function: trying to get the client 'into contact' with his feelings without him being overwhelmed by them.

2) In psychotherapy, one of the frequent goals is to show the client that much of his 'fear of falling' [Lowen] is imaginary and that the abyss is much farther away than he thinks. This, too, is a difficult and often fearful type of exercise.

> Human beings hardly ever control the "full trajectory."
> They control the variables that matter to them.

Yes. And bodily (and mental) integrity matters a great deal.

> "Stabilizing control" is something of a misnomer,
> suggesting that all that a control system does is to keep
> something constant. More generally, it makes the perceptual
> signal track the reference signal.

Exactly how would you know that the jumper follows a reference signal when for the very first time she jumps higher than she ever did before? How does the reference signal get established in the first place? I do not allow the answer that it is an 'imagined' reference signal; that would be impossible to either prove or refute and therefore unscientific [Popper]. I do allow the answer that the reference signal is discovered 'by accident', through trial and error learning. But that would mean that the very first time there was no reference that could be followed, i.e. that not _all_ behavior (here: peak performance) is control of perception.

Date: Mon Feb 08, 1993 9:05 am PST
Subject: book review from another list

[From: Bruce Nevin (Mon 93028 11:17:05)]

This book review looked like it might interest CSG folks. Disclaimer: as is often the case, I don't know this material and can't comment on the content. It looks useful and interesting.

Bruce bn@bbn.com

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

Neuron Digest Tuesday, 2 Feb 1993
Volume 11 : Issue 8

[many other items deleted, see list at end -- BN]

Subject: Wet brains as constraints on neural networks
From: Harry Jerison <IJC1HJJ@MVS.OAC.UCLA.EDU>
Date: Mon, 25 Jan 93 15:03:00 -0800

[. . .]
From some recent contributions in Neuron Digest, I sense an increasing interest in wet brains as constraints on neural networks. Last year I listed: V. Braitenberg and A. Schuez's ANATOMY OF THE CORTEX: STATISTICS AND GEOMETRY (New York, Berlin, Heidelberg: Springer Verlag, 1991. 249 pp. $39.00) as a useful book for theoreticians and mentioned that I had a review of it in press in CONTEMPORARY PSYCHOLOGY, a journal of reviews published by psychologists. With apologies to my AI and NN friends (see the concluding lines), and in the hope that they (and non-friends as well) will let me know if they are interested in this sort of communication, here are extracts (sans italics) from that review:

"Primarily determined from data on the mouse brain, (Braitenberg & Schuez) is a magnificent achievement in anatomical analysis and integration that deepens our understanding of the way a brain can work to generate a mind. That "mind" was described by Hebb (ORGANIZATION OF BEHAVIOR, New York, Wiley, 1949), and Braitenberg and Schuez provide the anatomical background for Hebb's cell assemblies and "Hebbian neurons" as living things rather than theoretical constructs or computer artifacts....

"The anatomy in Braitenberg & Schuez is a description of the geometry and statistics of the cortex (as promised in the subtitle) that would be required for understanding a brain as an information-processing system. We learn many new things. We learn something about the size of the brain as a system of connections. We learn, for example, that there are about 100 billion (10^11) synapses in the mouse's cortex and about 100 trillion (10^14) in humans. We learn that information-processing capacity is about the same PER UNIT VOLUME in brains of all mammals, and we are shown how these quantities are estimated. We learn that the majority of synapses in the mouse cortex are excitatory and not inhibitory. By comparing brains of altricial mice with those of precocial guinea pigs, . . .Braitenberg and Schuez (p.137) 'conclude (their) search for the anatomical traces of learning with two likely candidates, the number of vesicles on the presynaptic side of a synapse, and, for synapses involving spines, the thickness of the spine.'

"Traces of learning are about engrams, and the preponderance of excitatory synapses is a necessary element for a nervous system that would produce Hebbian cell assemblies. Braitenberg and Schuez also estimate the connectivity of neuronal systems, to suggest how cell assemblies would, in fact, be assembled.

"Why do psychologists need such information? They need it to write acceptable theories of the mind of the sort that Hebb wrote. Some years ago, I criticized Dalbir Bindra in these pages (CONTEMPORARY PSYCHOLOGY 22:417-419, 1977) for relying heavily on Hebb's model, which I thought too speculative to support Bindra's theorizing. Hebb himself came to Bindra's defense, criticizing me as seeing the brain "from afar," and implying that cell assemblies had already been demonstrated. I argued with Hebb (CONTEMPORARY PSYCHOLOGY 22:849-50, 1977), but after reading Braitenberg and Schuez I am less inclined to argue. The assemblies may remain theoretical, but they are close to the status of atoms as understood at the turn of the century, when the skeptical physicist-philosopher, Ernst Mach, would challenge other physicists to show him one. I am still put off by the artificial intelligencers, by 'neural networks' that exist only in computer programs, and by cognitive scientists who have become true believers in these silicon figments. Braitenberg and Schuez have begun to convert me."

(From CONTEMPORARY PSYCHOLOGY 1992, 37:927-928.)

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

Info about this digest:

> Today's Topics:
>   building energy predictor shootout -- data available by anon ftp
>   How to build a Boltzmann machine?
>   new cluster version available

> The Hunt for Info
> back prop nn refs request
> Is there a Fuzzy systems digest?
> Wet brains as constraints on neural networks
> Postdoc - computational neurobiology
> Job Offer: Research on Financial Analysis in Santa Fe NM
> Computational Biology Faculty Position

> Send submissions, questions, address maintenance, and requests for old
> issues to "neuron-request@cattell.psych.upenn.edu". The ftp archives are
> available from cattell.psych.upenn.edu (130.91.68.31). Back issues
> requested by mail will eventually be sent, but may take a while.

Date: Mon Feb 08, 1993 10:27 am PST
Subject: What's it all "about", Alephi

[From Rick Marken (930208.0900)] Oded Maler (930206.1000) --

>Have fun on the mountain, meanwhile we ask Rick to build us the golden veal

This is a wonderful biblical joke and quite appropriate since I love to worship golden veals and my grandpa looked just like E. G. Robinson (who played the guy who got the golden veal cult going in DeMille's "Ten Commandments"). So while Moses .. er Bill Powers... is away, I say let us worship the golden calf's of psychological science -- the cause-effect model and statistical significance tests.

>Int-a, "aboutness" is very
>important as I indicated above unless you adopt solipsism, deny the
>fact that your thoughts and perceptions are about something, and
>abandon the illusion of language as meaning carrier (rather than
>social ritual).

PCT is not solipsistic; the emphasis on perception is not aimed at promoting solipsism. It is done (I think) to emphsize the importance of modelling. The great achievement of physical science is that it created VERY accurate, predictive models of the cause of our perceptual experiences. These models are SO successful that we think of them as reality itself; but they are, at best, only possible perceptions that explain other perceptions we have (or will have if we operate on our perceptual world in certain ways).

Language is "about" our perceptions which are about an external reality that we know only in terms of our current models of that reality (the models of physics, chemistry, etc). The goal of PCT as a model (of another aspect of our perception -- our perception of purposive behavior) is to be, at least, consistent, with the best models we have of "external reality" in the places where the PCT model interfaces with those models. This attempt to be consistent with other, extremely successful models of "reality" is another thing that distinguishes PCT from most other models of living systems(to the extent that they are models). For example, one problem with "reinforcement" theory is that it is often formulated in a way that is inconsist with what we know of "external reality" from the models of physics and chemistry. Skinner, for example, thought that some objects were reinforcers as though reinforcement were a property of the object itself. But physics and chemistry reveal only the usual atomic components in both reinforcers and non-reinforcers. The PCT model of reinforcement (as a controlled percetual variable) does not require the invention of new entities "out there" that physics and chemistry don't need; reinforcement is just a physico- chemical variable in the external world; it becomes a perception in an organism (consistent with models of neurology); that perceptual signal is compared with a reference signal and an output (error) signal produces muscle actions that affect (via that external world) the reinforcer and, thus, the perception thereof, moving it toward the reference value. Everything in this model is compatible with our models of physics, chemistry, neurology and physiology. PCT actually puts us in closer touch with the best notions we have of what is "out there"; it is not solipsistic at all (unless, by solipsistic, you mean that we simply cannot "see" what is REALLY on the other side of our perceptual experience; PCT IS solipsistic in this sense; but this kind of solipsism is what justifies MODELLING; if we could somehow "see" the cause of our experience directly -- whatever that would mean -- then science would be unnecessary. I guess some religious types imagine that god is the cause of their experience and that through proper faith and worship they can "see" god. These folks don't need science at all -- a psychiatrist, maybe, but not a scientist).

Best Rick

Date: Mon Feb 08, 1993 1:22 pm PST
Subject: clinical strategies

from Ed Ford (930208:1355) Dennis Delprato (930208)

>..are you implying that one solution that follows from PCT is the
>teacher saying something, such as, "Billy and Harold, please
>refrain from talking"? I am interested in identifying concrete
>clinical implications of PCT-what might be called principles,
>strategies, and procedures.

I would not tell the children what to do, but rather ask the children what they are doing. If I have found anything out about dealing with others through my understanding of PCT, it has been that you can only access a living control system through questioning, not through telling. When you ask them a question, they deal with their world, and if you tell them what you think, they deal with you and what you're saying. By asking questions concerning their reference levels (what they want), their perceptions (what the rules are, what they are doing); evaluative questions that get them to compare what you've been talking about (is what you're doing getting you what you want? or is what you're doing against the rules?) and then test the strength of the reference signal (are you willing to work at resolving your problem?), you have a strategy for accessing their world and, more importantly, getting them to deal with their world responsibly.

Even getting children to make plans involves PCT. The measurable goal is the reference signal, the measurable feedback is the controlled variable, and the distance at any one time on a chart showing the goal and feedback is the perceptual error. The chart itself shows the historical progress of the child. This kind of PCT charting of the child's progress gets away from the vague planning that is generally found in schools, corrections, etc.

If you would post your present address, I'll send you several cards I give to teachers, parents, etc. who take my course as well as an outline of what I do. My strategy is outlined in my book, Freedom From Stress, chapter 9 and 10. The basis for what I do is Reality Therapy, but I have modified it to fit PCT.

Best, Ed
Date: Mon Feb 08, 1993 2:34 pm PST
Subject: Re: perception not solipsistic

[Martin Taylor 930208 16:30] (Rick Marken 930208.0900 to Oded Maler 930206.1000)

>>Int-a, "aboutness" is very
>>important as I indicated above unless you adopt solipsism, deny the
>>fact that your thoughts and perceptions are about something, and
>>abandon the illusion of language as meaning carrier (rather than
>>social ritual).
>
>PCT is not solipsistic; the emphasis on perception is not aimed at
>promoting solipsism. It is done (I think) to emphsize the importance
>of modelling.

I think that "to emphasize the importance of modelling" is a very weak (if possibly correct) reason. The main point (I think) is that there really IS nothing else that can be relevant to whatever a living thing is doing. Whatever may be out there, it affects the actions of living things only as they perceive (in the sense that PCT uses the word). One can argue that it affects living things only as they perceive, since even if the source of a perception is unclear (e.g. the nausea of radiation sickness), if the world had no effect on the perception, the only possible effect on a living thing would be to make it non-living (then it would not perceive).

One could imagine acts that are not part of the control of perception, such as those induced by electrical stimulation of muscles or neurons, but if the acts or their consequences are not perceived, they might as well not have happened. One's own unperceived acts may affect the world in ways that are perceived, but not related to one's own acts. They then have much the status of the nausea induced by the ondetected radiation.

It is to cover the ability to learn to control these "undetected cause" perceptions that Bill Powers has the reorganization hierarchy separate from the main perceptual control hierarchy. By random reorganization one might conceivably learn to keep away from the radioactive area, rather than to take anti-nausea pills. Local reorganization, based on "the tummy hurts--fix the tummy" could never do that.

As long-time CSG members know, I have a different way of approaching this same problem, that does not have a separate reorganization hierarchy. Either way, the result of considering only what the living organism can perceive is far from solipsistic. It is the ONLY self-consistent way of dealing with the interactions of living things with a world outside themselves.

Subjectivism is not solipsism. Einstein's view is self-consistent, where Newton's is not.

>PCT actually puts us in closer touch with the best notions we have of

>what is "out there"; it is not solipsistic at all (unless, by
>solipsistic, you mean that we simply cannot "see" what is
>REALLY on the other side of our perceptual experience; PCT IS
>solipsistic in this sense; but this kind of solipsism is what
>justifies MODELLING; if we could somehow "see" the cause of our
>experience directly -- whatever that would mean -- then science
>would be unnecessary.

Just so. Martin

Date: Mon Feb 08, 1993 3:19 pm PST
Subject: Re: Good Data/No Model

[Martin Taylor 930208 17:00] (Rick 930202.1100 resp. to Martin 930202 12:20)

>>No, I don't have a model of what goes on during a button choice.
>
>Then why did you decide to do this experiment and collect this data?

Note the tense of my non-claim.

>I suspect that there was a model lurking in the background; it was
>the ol' cause-effect model. I would venture to guess that the
>model was something like:
>
>stimulus information --> processing --> output response
>
>The linear relationship between response delay (the independent
>variable) and d'^2 (the dependent variable) presumably reveals
>something about the nature of the processing stage. The response
>delay was probably thought of as something that affects the
>amount of stimulus information or processing time available for
>producing the output. Is this about right?

Probably. I can't speak for Schouten, who obtained the data, but it would probably have been right for me, if I had done the experiment around that date.

>>But whatever
>>that model might be, it must incorporate the result that the information
>>relevant to the choice becomes available at a linear rate after some time
>>delay.

>Not quite, I think. The model just must behave in such a way that the
>d'^2 measures obtained from the model (just as they are from the
>subject) are linearly related to the response delay measures (again
>obtained from the model as they are from the subject).

I don't see the difference between these two statements. The d'^2 measure converts directly and linearly into the information measure. (Perhaps I do see a subtle difference, though; in my statement I should have included "becomes available *to the final control system that is responsible for the finger moving to one button or the other*..." With this amendment, I think the two are paraphrases.

>The problem with claiming that there is no model underlying experimental
>results is that it makes the experimental results themselves seem
>very important. But experimental results are only important (and meaning-
>ful) in terms of the underlying model about which they presumably
>provide evidence. The results per se are not particularly important.

I wouldn't go along with this. I agree with the last sentence, but it seems to contradict the first. If you read the Occam's razor note, you can see why. If all you have are data, then you have no way to reduce their description, and thus you have learned nothing that you can use in other circumstances. You have no linkage and no prediction. If you have a hypothesis (e.g. a model or even a verbal description), you may be able to make a considerably reduced description. The better the data, the more precise the hypothesis must be to provide a good reduction (and thus be Occam-good) and thereby permit prediction and the linkage of the model to other situations. The models of PCT tend to be quite precise, which is why PCT modellers demand good data.

>All observations (even Kepler's) are made in order to test a model;
>it's the model that's important, not the observations per se.

No, I can't go along with this, at least not with the word "All". Exploration is still useful. Eventually it is the model that is important, for sure. But the observations can either precede or follow its development. Usually there is a feedback loop of observation->model restructuring->observation

We, as mobile animals, seem to have developed senses with different missions. Some map (vision is a prime example, but haptic touch is another) to lay out the possibilities for behaviour, some are primarily for alerting (the visual periphey, audition or the skin senses) so that the appropriate perceptions that lead to survival coan be controlled from moment to moment, and some are primarily for perceptions that are being controlled (vision again, smell, taste). Any mapping sense is presumably suited for providing perceptions that can be controlled, which is why vision appears in both roles.

I think the mapping function is quite important. When you say "all observations are made to test a model" you deny the mapping function. Kepler was not originally testing whether planetary orbits are ellipses (they aren't, in detail). Tycho Brahe, whose observations Kepler used, certainly was not. And when Kepler or Brahe tried to fit them by parameterizing ever finer epicycles, it just got complex. The observations were there. The elliptical orbit model followed. Likewise, Schouten's data may well have been gathered to help parameterize a model, but he didn't think of the information gathering implications. That was mine, based on seeing that he had data that I might be able to use. At the time I didn't know about PCT. Now I do, so I realize that the original notion I had at the time is presumably not correct.

But the data remain. And any future model must accomodate them.

>The point of the PCT demos and experiments is that all these different
>little "findings" can be shown to reflect(Very PRECISELY) the same
>underlying process -- CONTROL OF PERCEPTION. The findings themselves are
>important only insofar as they test the underlying model; it's the model
>that's important.

Yes, the PCT demos DO illustrate the fact of control of perception (at least in my mind they do). But once one has demonstrated it, one gets very little further information from more demos that say the same thing. The problem becomes to discover which of an infinity of possible control system structures is performing the controlling. That, we find out by looking at situations in which the subject (and necessarily any good model) fail to control. So we find poor control but good correlation between the model and the subject.

>So what is the finding of a PRECISE linear relationship between response
>delay and d'^2 other than one of the many random observations we can make
>about human behavior? What does this finding tell us about the organization
>of the system that produces it?

Again, we come back to Occam. A linear relationship between d'^2 and time is a signature of a linear gain of information. That's a concept that need not be added to any reasonable model that includes a limited sensory system such as an optic nerve or an auditory critical band. It comes for free. If your model wants to deny that a channel that seems on the face of it likely to be able to pass information at a stable rate actually does so, and that evidence of a linear information gain is due to some other factor, then that model has to have some pretty good reductions of descriptive length to compensate. In other words, it has to explain something else outside the domain of this experiment in a way that the linear information gain does not, and at the same time mimic the results of this experiment.

All that the data show in themselves is that information from the event of turning on one of the two lights was used by some mechanism that directed the fingers to one of two buttons at a rate of roughly 140 bits/sec. It says nothing about the organization that permits this to happen. That's a question that should be addressed in modelling the control hierarchy.

To finish this posting where it began: "No, I don't have a model of what goes on during a button choice"

Martin
Date: Mon Feb 08, 1993 6:48 pm PST
Subject: data and models

[From i.n.kurtzer (930208..2010)] Rick Marken (data argument w/ Martin)

i fully agree with you that data without an EXPLICIT model is useless at best--all data must have an implicit model, however it is often in a untestable form (i.e. assumptions), else all would be chaos. data not only precedes and follows from explicit models, but data CEASES with new models as that data is put into a scheme that renders it as insignificant and misguided. thereby, much data obtained through models that did not acount for organisms as perceptual control systems is useless (i.e. S-R methodologies). and there is no obligation to account for the said data. when marken pointed out how control could be seen from several views (blind man paper) it followed that the other views were no longer tenable, assuming parsimony and unification of disparate "data" are scientific objectives, AND THE DATA that followed from the views were catalogues of meaningless--they lost it in the translation.

on another note, my new eclectic (@%^#!!!) perception teacher has us doing black box labs (ex. IHTT via manual reaction times) so i've begun my subtle dialectic. ha!

i.n.kurtzer

Date: Tue Feb 09, 1993 9:30 am PST
Subject: Re: control article

[From Rick Marken (930209.0800)] Hans Blom (930208) --

> Only in the simplest of experiments one
>variable can be considered to be controlled, if at all.

You missed my point; I was not suggesting Levine and Loeb find THE one controlled variable involved in jumping; I was suggesting that they were simply assuming (as you are in your answer) that they knew what variable(s) is (are) being controlled. I was simply suggesting that Levine and Loeb might try some version of the test for the controlled variable before trying to model behavior. There are many different versions of this test but the reason for the importance of testing for controlled variables (and not taking for granted that you can tell, just by looking, what variables are being controlled) is described in the "Behavior in the first degree" paper in my "Mind Readings" book; copies are still available.

>I would maintain that one of the most important of an organism's
>objectives is, at all times, not to seriously damage itself. But
>that cannot be formalized by control in the usual sense of the
>word, that is a perception following a reference signal.

Really? How does the organism avoid "seriously damaging itself" unless it can perceive the degree to which it is damaged? How can it not seriously damage itself unless it can do things to keep the perception of the degree to which it is damaged at an intended level (zero, for example)? How can you model this with anything other than control theory??

>The control system is operating under constraints, i.e. it must _stay
>away from_ certain experiences with a high probability of success
>(catastrophy theory studies such problems).

It "must stay away" -- but how is that "must" enforced? Do the "constraints" somehow keep the animal from producing actions that have the consequences that the animal must not experience? This would give the environment rather amazing animistic capabilities, it seems. I can see how bad consequences are avoided if the animal is a control system. But I don't see how this is done if the animal is a catastrophe theory. Could you explain the catastrophy theory model of this phenomenon? My guess is that catastrophe theory might DESCRIBE the phenomenon; but I doubt that it explains why animals so rarely experience the catastrophes you mention -- like muscles being torn from their attachments.

Best Rick

Date: Tue Feb 09, 1993 12:16 pm PST

Subject: Re: Intentionality

[From Chris Malcolm] Oded Maler writes:

> I have no intention to start a comp.ai.phil.-style of discussion
> but I can't resist trying to reformulate your suggestion.
>
> What is the semantics/denotation/aboutness of a reference signal
> (inside an individual)? It is the set of states-of-affairs in the
> world, such that when sensed by the individual will cause a zero error
> in the corresponding comparator.

I think I would agree with that.

> Still a problem with levels - should you attribute intentionality
> to a signal in a motor-neuron, or a thermostat?

I can see that "intentionality" is a minefield of misunderstanding here! First, note that the technical meaning of "intentionality" in philosophy of mind is NOT to do with intention in the sense of purpose (although, confusingly enough, there _happens_ to be a relationship). In philosophy of mind "intentionality" simply means the "aboutness" of a symbol, no more. It does _not_ derive from "intentional" meaning "on purpose, although it looks as though it does. It is in fact derived from the extension/intension distinction in linguistics (not spelling of "intension" here), and foolishly enough hasn't preserved the distinguishing "s" when adding "ality". By the way, most dictionaries don't know this.

So, do thermostats have intentionality in this (aboutness) sense? Well, if we choose to say that the state of the switch is a symbol meaning either "too hot" or "too cold", then it is certainly true to say that these are wired up appropriately to the environment, and when present in the "mind" of the thermostat certainly cause appropriate behaviour. That is an argument for thermostats having (the most elementary kind of) intentionality.

The argument against -- which I incline to -- says that this kind of intentionality is only appropriate in systems with a level of propostionally governed behaviour, i.e., which combine their symbols into collections of propositions, and perform some kind of reasoning with these. In essence, this argument says that intentionality is a feature of symbols in a language, and one (or two) words is not enough for a language.

On the other hand, I definitely argue that a thermostat (by which I mean the whole complex of thermostat, heating system, room, etc.) does have a purpose, is a goal-seeking device, and thus is intentional (meaning, has a purpose).

Thus I argue that thermostats have intention (purpose) but not intentionality (the aboutness of symbols in a language-like system).

I think I'm going to have to excise this technical philsophical use of the term "intentionality" from my writings. Either that, or start spelling it "intensionality" -- but would anyone in the US notice that? :-)

Chris Malcolm

Date: Tue Feb 09, 1993 12:37 pm PST
Subject: Re: intentionality?

[From Chris Malcolm] Francisco Arocha writes:

>In some philosophical writings, ever since Brentano, the term
>"intentionality" has come to refer, in again some philosophical
>discourse (and maybe now in AI too), to the "aboutness" that Bill P.
>mentioned recently. However, it seems to me that Brentano (and some
>philosophers) confuse intentionality (a psychological category) with
>reference (a semantic category).

There is no doubt that by "intentionality" Brentano meant "meaning as consciously meant", i.e., both reference (semantics) and (awareness) psychology. Later writers have usually restricted the term to its semantic meaning, but the fact that some of them (e.g. Searle) consider, like Brentano, that semantics can't exist without full conscious awareness of the human kind conflates the distinction.

Chris Malcolm

Date: Tue Feb 09, 1993 4:49 pm PST
Subject: intentionality

[Avery.Andrews 93010.1135]

Two remarks on `intentionality'. First is that although `PCT epistemology' is often expressed in a way that makes it seems rather solipsistic (I'm not implying that it is, just that something tends to make it look that way), it seems to me to be entirely compatible with the kind of realism in vogue at places like CSLI. Martin Taylor made the point very well with his `CEV's (Complex Environmental Variables). A percept assumes one form inside the organism, as the firing rate of a neuron, for example, and another outside, as a complex and perhaps rather subtle property of the environment. To keep the Gibsonians happy, one can say that this transformation of form is normally achieved by means of lawful transformations of energy, tho I think there a cases where more chaotic and error-prone processes get involved.

The second has to do with meaning. One of the puzzles of semantics is how a word like `gold' or `plutonium' can have a meaning that is in some sense independent of the concepts in the brains of most of the individual speakers of the language. E.g. none of us could right now recognize plutonium & distinguish it from other substances without killing ourselves (though some of us could probably figure out some way to do it, given time & access to the right kind of library), but there is a sense in which we know what it means, & can use this knowledge effectively (to vote for or against making it, deciding whether or not to give Greenpeace money to hinder its being shipped around the world, etc.). I'd suggest that the meaning exists in part by virtue of arrangements, in the society at large, for correcting `errors' in the usage of the word. E.g. the society as a whole can be seen as constituting a control system that controls for people applying the word `plutonium' to a certain kind of stuff.

Avery.Andrews@anu.edu.au

Date: Wed Feb 10, 1993 12:01 am PST
Subject: Back to Egypt!

[From Oded Maler (930209.1900 ET)] Rick Marken (930208.0900)

* E. G. Robinson (who played the guy who got the golden veal cult
* going in DeMille's "Ten Commandments"). So while Moses .. er Bill

I'm glad MGM took care of your classical education.. :->

* PCT is not solipsistic;

I didn't say so, I just indicated that the question of intentionality in the sense of aboutness is not a superficial question. It may be however orthogonal to "mainstream" PCT.

* The great achievement of physical science is that it created VERY
* accurate, predictive models of the cause of our perceptual experiences.

I don't get it. Physical science cannot define objects using "objective" coordinates. Maybe I don't understand what you mean.

* These models are SO successful that we think
* of them as reality itself; but they are, at best, only possible
* perceptions that explain other perceptions we have (or will have
* if we operate on our perceptual world in certain ways).

* [Rick Marken (930209.0800)] Re: Hans Blom

* Really? How does the organism avoid "seriously damaging itself"
* unless it can perceive the degree to which it is damaged?

Do you need the experience of a car accident in order to avoid accidents?

* How can it not seriously damage itself unless it can do things to keep the
* perception of the degree to which it is damaged at an intended level
* (zero, for example)? How can you model this with anything other than
* control theory??
*
* It "must stay away" -- but how is that "must" enforced? Do the
* "constraints" somehow keep the animal from producing actions
* that have the consequences that the animal must not experience?
* This would give the environment rather amazing animistic capabilities,
* it seems. I can see how bad consequences are avoided if the animal is
            ^^^^^^^^^
* a control system. But I don't see how this is done if the animal is
                       ^^^^^^^^^^^^^
* a catastrophe theory.

It's all (your) perception!

Best regars --Oded

Date: Wed Feb 10, 1993 3:45 am PST
Subject: Re: control article

[Hans Blom, 930210] Rick Marken (930209.0800)

>>I would maintain that one of the most important of an organism's
>>objectives is, at all times, not to seriously damage itself. But
>>that cannot be formalized by control in the usual sense of the
>>word, that is a perception following a reference signal.

>Really? How does the organism avoid "seriously damaging itself"
>unless it can perceive the degree to which it is damaged? How can
>it not seriously damage itself unless it can do things to keep the
>perception of the degree to which it is damaged at an intended level
>(zero, for example)? How can you model this with anything other than
>control theory??

Sometimes an organism just cannot 'perceive the degree to which it is damaged', because there is no 'degree'. You cannot fall off a cliff just a little. When walking along the cliff, you can err in one direc- tion (on the safe side), but not on the other.

>>The control system is operating under constraints, i.e. it must
>>_stay away from_ certain experiences with a high probability of
>>success (catastrophy theory studies such problems).

>It "must stay away" -- but how is that "must" enforced? Do the
>"constraints" somehow keep the animal from producing actions
>that have the consequences that the animal must not experience?

More formal, then. Consider a car driving with constant speed on a narrow road with cliffs on both sides. The weather is a bit gusty. We consider just the position of the car relative to the middle of the road. Call this variable x. Model x as a function of time; x depends on 1) the x of a moment ago, 2) the way you move the steering wheel; call this influence u, and 3) the wind and other random influences on x; call these e. The model of the car's position is then something like (if you take a difference equation rather than a differential equation):

x (t + T) = a * x (t) + b * u (t) + e (t) (1)

where e (t) is unknown but hopefully its statistics are known. An often made assumption is that e=0 on average, and that its standard deviation is constant and known. Catastrophy threatens when the absolute value of x becomes too large. This imposes limits:

| x (t) | <= xmax for all t (2)

Assume also that there is no observation noise: a noiseless observation of x is available at all times. The problem is clear now: find a control law for u (t) that obeys (1) and (2).

In linear quadratic control, the time integral of the square of the error is minimized. That allows an occasional large error, provided such large errors do not occur too often. Here the situation is different: even one error > xmax is not allowed, but otherwise you can swerve all you like. There is no reference signal in the strict sense, although the control law will show that the average position will be the middle of the road.

Hope this makes things clear. No magical environment, just a different control law.

Best, Hans

Date: Wed Feb 10, 1993 6:10 am PST
Subject: connectionist paper: HRN memory

[From: Bruce Nevin (Wed 930210 08:53:54)]

The following is cross-posted from another list because it seemed that it might be of interest to at least some of us here.

Disclaimer: it's entirely possible that I don't know what they're talking about. Even if I do, ideas and proposals here are theirs, not mine, and it's not likely I can explain or defend them. What you see is what I got.

Bruce bn@bbn.com

-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-
To: neural-people@BBN.COM, machine-learning@BBN.COM
Subject: [mav@cs.uq.oz.au: Learning in Memory Technical Report]
From: aboulang@BBN.COM
Date: Tue, 9 Feb 93 16:53:14 EST

To: Connectionists@cs.cmu.edu
Subject: Learning in Memory Technical Report
Date: Mon, 08 Feb 93 10:33:21 +1000
From: Simon Dennis <mav@cs.uq.oz.au>

The following technical report is available for anonymous ftp.

TITLE: Integrating Learning into Models of Human Memory:
The Hebbian Recurrent Network

AUTHORS: Simon Dennis and Janet Wiles

ABSTRACT:

We develop an interactive model of human memory called the Hebbian Recurrent Network (HRN) which integrates work in the mathematical modeling of memory with that in error correcting connectionist networks. It incorporates the matrix model (Pike, 1984; Humphreys, Bain & Pike, 1989) into the Simple Recurrent Network (SRN, Elman, 1989). The result is an architecture which has the desirable memory characteristics of the matrix model such as low interference and massive generalization but which is able to learn appropriate encodings for items, decision criteria and the control functions of memory which have traditionally been chosen a priori in the mathematical memory literature. Simulations demonstrate that the HRN is well suited to a recognition task inspired by typical memory paradigms. When compared against the SRN the HRN is

able to learn longer lists, generalizes from smaller training sets, and is not degraded significantly by increasing the vocabulary size.

Please mail correspondence to mav@cs.uq.oz.au

Date: Wed Feb 10, 1993 7:22 am PST
Subject: PCT Personnel Management

[FROM: Dennis Delprato (930210)]

Recently met 2-person team who do full-time consulting and so on in the area of executive development, employee relations, and the like. About the first thing they said was that anyone working with people must always take into account that people self-regulate. I said that I knew something they may find interesting. Can anyone write or refer me to a brief paragraph that clearly ties PCT to the above work and supply 1-2 references?

Dennis Delprato psy_delprato@emunix.emich.edu

Date: Wed Feb 10, 1993 10:44 am PST
Subject: Back to Egypt!, re:control article

[From Rick Marken (930210.0800)] Oded Maler (930209.1900 ET)--

>I'm glad MGM took care of your classical education.. :->

Actually, MGM handled my biblical education; Paramount handled my classical education (Pride and Prejudice) and Public Television is currently handling my science education.

I said:

> The great achievement of physical science is that it
> created VERY accurate, predictive models of the cause of our
> perceptual experiences.

Oded says:

>I don't get it. Physical science cannot define objects using
>"objective" coordinates. Maybe I don't understand what you mean.

I meant that science comes up with models (Newton's laws, Dalton's atoms) that make it possible to predict aspects of what we observe (phenomena such as the future positions of the planets and the gaseous result of heating 1 liter of water until there is no more water) with great accuracy. I view the model entities (masses, forces, atoms, etc) and the simple laws describing how they interact as guesses about what causes aspects of our experience (like falling objects, exploding ballons, etc). I think some of the model entities can be considered "potential" perceptions because there are ways to create perceptions (if the model is right) that impress people as being more direct confirmation of the existence of the model entity. Thus, I have seen electron micrograph pictures of atoms (looking like little spheres all packed together); such pictures don't "prove" atoms exist any more than do the precise results of standard chemical experiments. They are just another confirmation of an aspect of the atomic model -- a perception that is the predicted result of certain operations (the operations of the electron microscope) on a material. But somehow the picture of atoms seems like more direct evidence of their existence than, say, the breakdown of water into two parts hydrogen and one part oxygen gas. But both the pictures and the volumes of gases are perceptions; the atoms are make-believe; guesses about what causes these perceptions to behave as they do.

I use the word perception to describe everything I experience; the desk in front of me, the voices in the other bay, the pictures produced by the electron microscope, the electron microscope itself, the moon, the stars, etc. All there is is perception; perception IS reality. The idea that there is some unperceived reality that underlies the reality that is direct experience (perception) is a rather sophisticated notion, I think, and it is the basis of both religion (animism), science (anti-animism) and PCT (animism and anti-animism applied appropriately).

I said:

>Really? How does the organism avoid "seriously damaging itself"
>unless it can perceive the degree to which it is damaged?

And Oded says:

>Do you need the experience of a car accident in order to avoid accidents?

In the same vain, Hans Blom (930210) adds:

>Sometimes an organism just cannot 'perceive the degree to which it is
>damaged', because there is no 'degree'. You cannot fall off a cliff
>just a little. When walking along the cliff, you can err in one direc-
>tion (on the safe side), but not on the other.

Hans claimed that organisms have the OBJECTIVE of not seriously damaging themselves and that this fact cannot be formalized by control in the usual sense. My point was that this objective cannot be achieved unless it can be perceived (and, hence, controlled). Oded and Hans then mention that we avoid car accidents and falling off cliffs even though we do not perceive them (or, when we do, it's certainly too late to achieve the objective of not seriously damaging ourselves). In fact, we rarely get in accidents or fall off cliffs because we do control variables that prevent these uncontrollable results. When you are driving, you don't control "the degree to which you are in an accident"; you control how close you are to other cars, the programs you carry out while driving (the rules of the road), your speed relative to bumps and wet spots on the road, etc,etc. The result of controlling all these variables is USUALLY no accident. The same is true of cliffs. You control your distance from the edge, your center of gravity when you are near the edge, etc. So in these cases we are not directly controlling the degree to which we are damaged. We are controlling other variables IN THE HOPE that by doing so we will not be damaged. But I think there are instances when we do explicitly control variables that could be called "the degree to which we are damaged" -- such as when we pierce our ears or get tattooed or cut off our foreskins or bind our feet. Although these things are often involuntarily imposed on other people, the people doing the imposing are controlling a variable that could be called "level of damage" and they want to control it at a level that's (qualitatively speaking) greater than "no damage" but considerably less than "fatally damaged".

I still vigorously disagree with Hans'claim that it is possible for organisms to have the OBJECTIVE of not seriously damaging themselves and that this fact cannot be formalized by control in the usual sense. What good is an objective if you can't control it (achieve the objective in the context of unpredictable disturbances)? and how can you control something if you can't perceive it? Perhaps I'm not understanding what Hans meant by "control in the usual sense". In what UNUSUAL sense of control do organisms achive the objective of not seriously damaging themselves? If this objective is achieved as a side effect of controlling other variables then it is not an OBJECTIVE (controlled result) at all; it is a lucky side effect. In fact, I think we do control variables that are "intrinsically" related to "damage"; these are the intrinsic, physiological variables that are maintained (indirectly) by the perceptual control hierarchy (which is reorganized when these intrinsic variables are not being held at their genetically determined reference levels).

I said:

>It "must stay away" -- but how is that "must" enforced? Do the
>"constraints" somehow keep the animal from producing actions
>that have the consequences that the animal must not experience?

Hans Blom (930210) replies

>Catastrophy threatens when the absolute value
>of x becomes too large. This imposes limits:

| x (t) | <= xmax for all t (2)

>The problem is clear now: find a
>control law for u(t) that obeys (1) and (2).

OK. This describes the problem.

>In linear quadratic control, the time integral of the square of the
>error is minimized. That allows an occasional large error, provided
>such large errors do not occur too often. Here the situation is dif-
>ferent: even one error > xmax is not allowed, but otherwise you can
>swerve all you like.

So you are proposing a particular kind of control model that satisfies this constraint. But this still doesn't mean that this is the way a person actually behaves in such a situation; this analytic approach to control system design may be fine in engineering contexts but I don't see what it can contribute to our understanding of living control systems. There are many other ways to model the control of x in your example; the "right" way must be determined by testing the model against real be- havior; not against catastophe theory. Most important, we don't even know that the driver is controlling the variable, x, that is controlled by your model.

>There is no reference signal in the strict sense,
>although the control law will show that the average position will be
>the middle of the road.

There is no reference signal because you didn't put one in explicitly and you set up the model so that the constant of integration (zero) corresponds to the value of x that is defined as "center of lane". In fact, when I drive down a road (mountains on both sides or not) I am perfectlly capable of changing my reference for the position of the car; in fact, I have even pulled off onto the side on one-lane mountain roads; I did this, not to test PCT but to avoid fast moving trucks that were proceeding directly toward me; nevertheless, this action accomplished both goals (avoiding the truck and testing PCT).

Best Rick

Date: Thu Feb 11, 1993 3:50 am PST
Subject: Re: re: control article

[Hans Blom, 930211] (Rick Marken (930210.0800))

>Hans claimed that organisms have the OBJECTIVE of not seriously
>damaging themselves and that this fact cannot be formalized by
>control in the usual sense. My point was that this objective
>cannot be achieved unless it can be perceived (and, hence, controlled).
Perception is not the only human capability that we depend on to control our behavior. Sometimes memory will do: a child will stay away from a hot stove after having been 'bit' by it only once. Sometimes 'knowledge', such as from a newspaper, will do: stay away from Chernobyl for a while. In neither case do you control for an exact distance from the feared location, you just want to keep at least a minimum distance away from it.

Maybe we have a different conception of what perception is. For me, perception is everything that my senses register and what can be derived from that. You might include memory as some type of 'observation' through 'inner senses'. Is that what you mean?

That leaves the discrepancy of wanting something and not wanting some- thing. More philosophically, I think that this distinction explains what gives us freedom. There is not one optimal location that is dictated by a match between our inner drives (reference levels) and our perceptions of the outside world. I do not dispute that we have reference levels and that we use our perceptions to get us close to them. I just want to add something like 'negative reference levels', things to stay away from. Freedom is a name for ranges in N-dimensional objective space where you can move about 'at will', because the objective function is flat. It is as if you try to find the highest peak in a mountain range and once you get there you discover a wide, high altitude tableland.

An example: you get conflict when the heater is set to 22 degrees (we use Centigrade) and the airconditioner to 20 degrees. You get a region of 'freedom' if the heater is set to 20 degrees and the airconditioner to 22 degrees.

> When you are driving, you don't control
>"the degree to which you are in an accident"; you control how close
>you are to other cars, the programs you carry out while driving (the
>rules of the road), your speed relative to bumps and wet spots on
>the road, etc,etc. The result of controlling all these variables
>is USUALLY no accident. ... We are controlling other variables

>IN THE HOPE that by doing so we will not be damaged.

As a control systems designer, I must seriously object. We do not create control systems 'in the hope that' they function correctly; hope has no place in the model. We do not rely on things going right only 'usually'. We specify an objective function that we know will lead to a correct design. And if we cannot guarantee correctness, we will at least strive for optimality in some sense, such as longest mean time between failures or longest time before first failure (hints of catastrophe theory again). Would evolution be sloppier, given its billions of years of experiment- ation?

> Perhaps I'm not understanding
>what Hans meant by "control in the usual sense". In what
>UNUSUAL sense of control do organisms achive the objective of
>not seriously damaging themselves?

By 'control in the usual sense' I meant the type of control that is discussed mostly in CSG-L, the type that all Bill's models are based on. In the engineering literature, it is called by different names, such as linear quadratic control or PID control. It assumes that the plant to be controlled is linear (or that its non-linearities can be neglected), it relies on linear relations in the controller itself, and its objective function can be shown to be the average of the minimum of the square of the deviation between a setpoint (reference level) and an observation. This is where the notion of a (possibly time-varying) setpoint comes in. In the multi-dimensional case, the objective function is the average weighted sum of the minima of the squares of the deviations, where the weights must be specified a priori. Weights introduce the notion of (relative) _importance_.

'Unusual' types of control exist in great variety; the system to be controlled can have pronounced non-linearities, the relations in the controller can be non-linear or the objective function can be something different from a time-averaged square of deviations. Over the last forty years or so a fine theory for the design of linear control systems has come into being. We do not as yet, however, have a theory to speak of that allows a systematic design of non-linear control systems, and I think that there will never be one. There just does not seem to be a _systematic_ approach. The general opinion is that you need to know the major characteristics of the system to be controlled in order to be able to design a well-behaving controller for it. In particular, control has been shown to be most difficult to design if the system to be controlled has (non-accessible) memory states. Very few control engineers work on non-linear problems. But some do.

Constrained control, as in my example, can be approached (there are also other ways) with an objective function that is the average weighted sum of the minima of the Nth power of the scaled deviations, with N even and large. The minimum is of course with respect to the control action:

$$J = \min_{u(t)} E \left\{ \sum_{t=0}^{t=tmax} (x(t) / xmax)^{1000} \right\}$$

For | x (t) | < xmax, the contribution to the objective function is negligible, for | x (t) | > xmax very large. That is exactly the objective.

>So you are proposing a particular kind of control model that satisfies
>this constraint. But this still doesn't mean that this is the way a
>person actually behaves in such a situation; this analytic approach to
>control system design may be fine in engineering contexts but I don't
>see what it can contribute to our understanding of living control systems.

I assume that evolution, through a harsh billion year long struggle for survival, may have come up with some pretty clever solutions to the control problems that have arisen. E. coli has a funny (partly random) but clever control law that results in what is called a biased random walk. This 'primitive' control law serves it quite well; coli is far more numerous than homo sapiens. Higher organisms have other (better?) control laws, some of which we seem to have more or less uncovered (control of voluntary muscles in humans) and which resemble linear quadratic control, at least as long as muscles function well within their force limits. Linear quadratic control works well in stabilization, i.e. stand-still and slow movements. In other cases, we know that there are better control laws. An example of that is when peak performance is required and the forces that muscles can deliver come to their limits. In that case, the non-linearities of the actuators cannot be neglected anymore and linear quadratic control becomes sub-optimal. Intuitively I agree with Bill Powers when he supposes that there is only one control law that governs the control of muscles. Linear quadratic control is, in my opinion, its more readily understandable 'special case', just like Newtonian physics is a more readily understandable special case of general relativity.

>There are many other ways to model the control of x in your example;
>the "right" way must be determined by testing the model against real be-
>havior; not against catastophe theory. Most important, we don't even
>know that the driver is controlling the variable, x, that is controlled
>by your model.

We do not know that the driver wants to keep on the road rather than fall off the cliffs? Then what _do_ we know?

>There is no reference signal because you didn't put one in explicitly

Right.

>and you set up the model so that the constant of integration (zero)
>corresponds to the value of x that is defined as "center of lane".

Wrong. I could have specified the limits as follows:

x (t) > xmin
x (t) < xmax

Such a coordinate change would not make any difference for the resulting control.

What is so difficult in accepting that there are some things that we want and other things that we do not want? If a model cannot handle negation, too bad for the model.

Best, Hans

Date: Thu Feb 11, 1993 10:41 am PST
Subject: Re: re: control article

[Martin Taylor 930211 10:00] (Hans Blom, 930211)

Never having met Bill Powers, I don't know whether he has any hair to pull out after reading, the following, but I'm willing to bet he would if he could:

>By 'control in the usual sense' I meant the type of control that is
>discussed mostly in CSG-L, the type that all Bill's models are based on.
>In the engineering literature, it is called by different names, such as
>linear quadratic control or PID control. It assumes that the plant to be
>controlled is linear (or that its non-linearities can be neglected), it
>relies on linear relations in the controller itself, and its objective
>function can be shown to be the average of the minimum of the square of
>the deviation between a setpoint (reference level) and an observation.

Since Bill is away for another week, it remains for the followers to respond. One of Bill's points repeated over and over is that there is NO requirement for, or consideration of, linearity in the control loop--ANYWHERE. There is a problem with non-monotonic relationships, which has been addressed in several exchanges, but monotonicity is the strongest requirement put on what one might call "control in the normal Bill Powers sense." Bill has demonstrated this with many kinds of nonlinearity. Rick's spreadsheet demonstrates it where the nonlinearity goes so far as logical statements (e.g. a<b is true). All the comparator need do is determine the sign of the error. Anything else may be used, and may affect the dynamics of control, but emphatically, linearity is NOT expected.

As for avoidance issues, and what constitutes "perception," were you reading CSG-L during the discussion of alerting systems and attention deployment last year? If not, you can probably find it in the archives. We have indeed talked about ECSs with dead zones, in which the error is effectively zero over some range of (perception-reference) (this can be done in the comparator or in the output gain function). Your car not going over the edge can readily be handled this way, although a function more like a cubic would seem more useful than a threshold function.

Perception is whatever comes out of the perceptual input function of an ECS., What goes into that function can be direct sensory input (unusual), the perceptions of other (lower) ECSs, or the results of imagination, which may depend on memory or could be the results of output from some other ECS (or its own output, very probably). The latter is an aspect of planning, as it is usually considered.

When Rick talks about the organism not controlling a perception of survival, and things like that, the point is that survival is not in itself a perception that can be under control. But surrogates can be, and organisms that did not happen to control those surrogate perceptions very well do not have many descendants. We control for not going too recklessly near cliff edges because we have anscestors who survived. It is only in very recent times that we have had the ability to TELL children that it is dangerous. But we don't need to. Infants won't crawl over the visual edge of a drop even when the drop is actually covered by a strong glass sheet. They aren't controlling for "not getting hurt." They are controlling for something they can see--proximity to the edge, I would presume.

The fundamental point of PCT is "if you can't perceive it, you can't do anything about it." Everything you do is related to something you can perceive. Everything you do may have side-effects that you can't or don't perceive, but that is of interest only to someone else, not to you, even though it might kill you in the end.

And yes, references for avoidance are easily accomodated in the theory. A saturating nonlinearity in an ECS with a positive feedback loop will do it.

Martin

Date: Thu Feb 11, 1993 12:25 pm PST
Subject: Re: control

[From Rick Marken (930211.0800)] Hans Blom (930211) --

I will try to answer your points but first I must ask if you have ever read Powers' books ("Behavior: The control of perception" and "Living control systems I & II")? I don't mean to be presumptuous, especially since you are apparently a control engineer, but it seems to me that some of your points attribute properties to the PCT model of living systems that it just does not possess. For example, you say:

>By 'control in the usual sense' I meant the type of control that is
>discussed mostly in CSG-L, the type that all Bill's models are based on.
>In the engineering literature, it is called by different names, such as
>linear quadratic control or PID control. It assumes that the plant to be
>controlled is linear (or that its non-linearities can be neglected),

Have you seen the ARM demo? Plenty of non-linearities in that "plant" (the muscles, environment, etc). Have you seen Bill's 1978 Psych Review paper (reprinted in LCS I) where (among other gorgeous things) he shows how people (and the PCT model) can handle a CUBIC "plant" function (about as non-linear as you can get) when controlling a cursor (the function relates handle "outputs" to cursor "inputs". One of the "raisons d'etre" of PCT is that people behave (produce consistent results) in a highly non-linear environment (plant); that, according to PCT, is one of the benefits of the negative feedback organization; it obviates the need for precisely computed outputs that are the inverse of the non-linear effects they have on controlled results; organisms control PERCEPTION, not output. The non-linearities are one reason why anything other than a control organization never evolved.

>Perception is not the only human capability that we depend on to control
>our behavior.

Perception does not control our behavior; perception is CONTROLLED by actions; controlled perception IS behavior.

>Sometimes memory will do:

Please, show me the diagram of this model; the only sense I can make of it is that references can be set based on remembered perceptions. This is all part of PCT; try the chapter in "Behavior: The control of perception" on MEMORY.

>a child will stay away from a hot stove after having been 'bit' by it only
>once.

How does it "stay away"; my guess is that after the experience of being burned the child starts controlling it's distance (a PERCEPTION) from the stove differently -- it set's it's reference for distance at a value that might be called "far" rather than "close". It's still control of PERCEPTION, not memory.

People can control memory, by the way. It's called "imagination" (the model of this process is discussed in the Memory chapter in BCP). The problem (as many of us daydreamers eventually learn) is that imagination doesn't help us get things done. I can imagine that I'm outside watching the house burn down. But that will be the last thing I imagine unless I get my butt out of the house. The latter involves controlling the perception of my butt; so that I perceive it (not imagine it) outside of the burning house; this is CONTROL OF PERCEPTION -- the USUAL kind.

>I do not dispute that we have reference levels and
>that we use our perceptions to get us close to them. I just want to add
>something like 'negative reference levels', things to stay away from.

We don't "use our perceptions" to get close to reference levels; we use the effects we produce on the environment (linear or not) to influence our perceptions to keep those perceptions tracking internally specified reference levels. 'Negative reference levels' are just settings for perceptions that we give names to like "avoidance".

>As a control systems designer, I must seriously object. We do not create
>control systems 'in the hope that' they function correctly; hope has no
>place in the model. We do not rely on things going right only 'usually'.

You can produce a control system that controls, say, it's internal termperature perfectly, compensating for normal and even extreme disturbances perfectly. Then the system is hit by a meteorite; no more control. This is apparently what happened to the dinosaurs; perfectly designed control systems; bad luck with the asteroid. That's what I meant in the driving example. You can control all the "driving" variables perfectly -- and then get hit by a drunk coming out of nowhere at 120 mph. Shit happens. The fact that it does says nothing negative about the designer of the control system (be it control engineer, evolution or god).

>What is so difficult in accepting that there are some things that we want
>and other things that we do not want? If a model cannot handle negation,
>too bad for the model.

Isn't "not wanting" just a way of saying that you want something at a particular level -- like zero (not wanting to taste lima beans -- ie. wanting zero amount of that taste perception) or 1000 (like not wanting to be near a lawyer -- ie. wanting to be 1000 miles away from the nearest lawyer). "Not wanting" can also mean not caring one way or the other; PCT has this too; these are uncontrolled perceptions.

So, I certainly "accept" that people want things and don't want other things; but these is just words. Both "wanting" and "not wanting" are examples of the same underlying phenomenon -- CONTROL.

Best Rick

Date: Thu Feb 11, 1993 1:04 pm PST
Subject: Branching out

[From Rick Marken (930211.1200)]

I sent the following note to Gary Cziko (the God of CSG-L) suggesting the possibility of advertising CSG-L to people who subscribe to Psycholoquy -- the psychology bulletin board (or whatever the hell these things are called). He liked the idea but suggested that it would be best to get a "sense of CSG-L" before doing something that might get a whole tribe of people posting to CSG-L who know zilch about PCT. So here is a copy of my post to Gary; I think it would be fair to wait a couple of weeks for the CSG-L opinions on this before doing anything rash.

Hi Gary! A thought just occurred to me.

I just got my next door office buddy onto Internet and showed him how to subscribe to Psycholoquy (I didn't show him how to get onto CSGNet because 1) I want to remain friends with him and 2) I don't want him to know how I spend most of my time at work). He made a print out of the Psycholoquy subscriber list and there are a bazillion people on it. So it got me to thinkin: even though Psycholoquy will not publish our fringe (read -- "elegant") CSG material, perhaps we can at least start a dialogue with an open-minded subset of the Psycholoquy readership. We could do this by placing an "ad" for CSGlist in Psycholoquy. I think you should be the one to do it because 1) you "own" the list 2) you are a real academic psychologist and 3) you seem to be a tad more diplomatic than I.

Why not send a note to Harnad at

harnad@pucc
harnad@clarity
harnad@psycho.princeton.edu

(there are more but I'm sure one of these will work). Say something like:

"I would like to put a notice in Psycholoquy indicating the availability of an Internet discussion group on purposeful behavior. The group, csg-l, discusses the implications of a control model of purposeful behavior for studies of language, information processing, AI, chaotic attractors, social behavior, neural networks, psychopathology, evolution, operant behavior, consciousness, etc. You can subscribe to csg-l by ..."

Well, you get the idea. You should say it in your own words, of course, but that's my idea. Do you think it's reasonable? There are about 6500 names on the Psycholoquy directory and there are probably even more who are unlisted (I think I recall hearing that there were 20,000 subscribers!!!). Many are "well known" psychologists and I'm sure there are lots of "impressionable" grad students too. While I don't want the quality of csg-l to go down, I think an injection of new blood might breed some new and exciting directions for discussion and research. And it would sure be nice to talk directly with the "enemy".

Best Rick

Please reply to CSG-L, not to Gary or me. Thanks Rick

Date: Thu Feb 11, 1993 3:20 pm PST
Subject: Re: Branching out

[Avery Andrews 930212.0906] Rick Marken (930211.1200)

My initial reaction is that it would be better to first write something that gets accepted by psycholoquy, and use *that* as the advertisement. (I'm still working on my piece, but at a reduced rate - I should be able to post the next installment pretty soon).

On a somewhat different note, I wonder if `feedback' isn't actually a dirty word (two four-letter ones). The problem with it is that it suggests that the organism's perceptual functions are able to draw a distinction between the effects of what the organism does, versus the effects of the disturbances, hence terms such as `proprioception', `exterorecption', and `exproprioception'. I'm getting the impression in some of my reading that this confusion actually exists & causes a certain amount of trouble.

Avery.Andrews@anu.edu.au

Date: Thu Feb 11, 1993 3:35 pm PST
Subject: Re: Branching out

[Rick Marken (930211.1430)] Avery Andrews (930212.0906)

>My initial reaction is that it would be better to first write something
>that gets accepted by psycholoquy, and use *that* as the advertisement.

Sounds good; but at least three of us (Bill P., Tom B and myself) have independently tried MANY TIMES over the last 12 years to get papers through the gates kept by Harnad (Psycholoquy, BBS). My proposal was based on the assumption that getting PCT stuff in front of this audience seems (to me) hopeless. If you think you can do it, great! Good luck.

Rather than play this acceptance game, I think we should just let it be known, to those who might care, that we exist. Then they can read our stuff without the evaluative benefit of the referees at Psycholoquy -- and make up their own minds.

Best Rick

Date: Thu Feb 11, 1993 3:44 pm PST
Subject: Re: Branching out

[Martin Taylor 930211 18:00] (Rick Marken 930211.1200)

If PSYCOLOQUY (I'm sure that's spelled wrong) would accept a CSG-L invitation to subscribe, it seems like a good idea. But I would change Rick's list of implication areas to "for studies of all areas of individual psychology, from sensory-motor to clinical, from psycholinguistic to connectionism" or something like that. After all, people who try naively to talk about "classical" approaches through AI, chaos, and so forth tend to get jumped on quite vigorously in CSG-L.

The core of PCT is what Rick keeps reiterating under the two guises: "It's all perception" and "It's all control." Those statements apply to all the areas of psychology. AI and chaotic attractors are fields of study in their own right, but when they get applied to psychology, they don't mesh easily with PCT, and it might be misleading to advertise that CSG-L is a forum for discussing them.

Otherwise, it's a good idea.

Martin

PS. I may think differently tomorrow! It happens.

Date: Thu Feb 11, 1993 3:59 pm PST
Subject: Re: re: control article

[Avery Andrews 920212.1024] (Hans Blom, 930211, & follow ons)

I wonder of some of the friction between control system engineering and PCT might be due to the fundamentally different aims of the two subjects. A control system engineer presumably needs to be sure that their design will work properly, before advising anybody to spend money to build it. It is therefore a high priority to find a convenient subset of the possible control systems that have nice mathematical properties that allow things to be proved about them.

PCT on the other hand is concerned with identifying the control systems that natural selection and learning methods have come up with, and these operate completely differently from the way engineers do - they don't care at all about proofs and nice mathematics, but simply experiment massively in all directions, and cull the variants that don't work.

Bill's theory doesn't assume linearity, not because he has any mathematical methods for saying anything interesting about nonlinearities, but because proving things about control systems is just not a priority for PCT, at least at its present rather primitive stage of development. Rather figuring out what kinds of systems are arguably present in living creatures is the main task.

The Arm parameters, for example, were, I believe, settled upon mostly by trial and error.

Date: Thu Feb 11, 1993 10:54 pm PST
Subject: Re: control

[from Ray Jackson (930211.1145)]

for Rick Marken, 930211, and Martin Taylor, 930211, also Kent McClelland:

regarding Hans Blom (930211) --

Rick & Martin:
I really appreciate the lucid remarks to answer Hans' concerns. Some of us feel we have a good handle on PCT basics but, in these cases, commentary such as yours on the foundations of the theory helps to reinforce and further define the things we already know.

Kent:
Congratulations! Your PCT & Sociological Theory paper is a remarkable and diligent work. Gary was kind enough to download it to me and I'm still digesting it. A tremendous effort which will have a significant impact as a resource for all of us trying to make PCT work in the real world. Thanks.

Regards, Ray

P.S. Martin, you ought to be careful who you are associated with...

Date: Fri Feb 12, 1993 12:23 am PST
Subject: fowler and turvey 1978

[Avery Andrews 930212.1918]

I haven't had time to really think about this, so maybe I'm going off half-cocked, but it seems to me that one of the PCT papers that ought to be is a little discussion of the Fowler & Turvey 1978 piece, along the following:

They are correct in noting that the basic model of PCT (without reorganization) assumes monotonicity (of effector efforts w.r.t. changes in perceptual effort), and it is also probably underestimated the importance of non-monotonic situations. But they are wrong in claiming that the theory as a whole cannot account for people's ability with their experiment, since the proposals about reorganization form the basis of a straightforward solution, which Rick has already implemented, and demonstrated to mimic the behavior of actual subjects (Mind Readings 22-23).

This paper might be a triviality, in the sense that from the point of view mere content, Rick has already said it all, but I think it would be worth getting it written out in full and agonizaing detail.

Avery.Andrews@anu.edu.au

Date: Fri Feb 12, 1993 2:44 am PST
Subject: Re: Branching out

[Oded Maler 930211.1100-ET?] Avery 930212.0906 Rick (930211.1200)

* My initial reaction is that it would be better to first write something
* that gets accepted by psycholoquy, and use *that* as the advertisement.
* (I'm still working on my piece, but at a reduced rate - I should be able
* to post the next installment pretty soon).

I agree. I think that a paper entitled "An alternative approach to motor-control" *can* be written and accepted to BBS. It should be written around the lines of the FB too slow discussion. It need not contain the slogan "B is the C of P" and the pecentage of "internal" PCT references should not exceed 10%. It need not speak of a Newtonian revolution and all that. All these precautions are needed to convince the referees (and some potential co-authors as well :-) that this is not a bunch of lunatics.

I would suggest that the paper will be developed along the following lines:

1) Definition of the domain - low-level snesory-motor behavior (up tp the level of pointing-like behavior).

2) Discussion of ther alternatives for control schemes. Suggesting the PCT model. Objective vs. perceptual systems of coordinates. Closed-loop vs. input->state->output.

3) Discussion of the limited significance of psychophysical experiments. The significance of neurological and phisiological findings.

4) The difference between control-engineering/robotics criteria and PCT criteria for the adequacy of models.

5) Description of Little Man and its signficance (implicit approximate computation of inerse kinematics). Limitations. The replicated tracking experiments.

6) Hierarchy and time-scales.

7) CPGs and motor "programs" from PCT's point of view.

8) SPECULATIONS about higher levels and the problems they pose.

I think the putting some deadline of, say, 6 months for writing such a collaborative paper with contributions from many participants, will lead to the work being done. If after all the paper will be rejected from BBS, it will be already in a form of a book which can be published either by the CSG or (preferably) by a more established house - e.g. MIT Press (they will publish everything).

I think the paper should take advantage of the knowledge of Greg and recently Avery of the other approaches to motor control. The criticism of Rick and Tom on traditional experiments and observability in psychology. The mathematical background of people knowledgable in control engineering, info theory and robotics (I stopped mentioning names)... etc. (and knowledge of native English speakers:-) And not least importantly, of people who had the experience of being "mainstream" in some sub-stream of established science, and will have the necessary sensitivity for not presenting the paper as a complete outsider.

--Oded

Date: Fri Feb 12, 1993 8:00 am PST
Subject: 12 SHORT TOPICS - RKC

[From Bob Clark (930212.10:30 am EST)]

Oded Maler (930203), Rick Marken (930202.1200), Martin Taylor (930202.1345)
Bill Powers (930202.0138), others

ANTICIPATION
"Suitcase Flinging" concerns "anticipation." This word (anticipation) has been used in this connection, but without being really tied to HFB Theory very well. A common example: the time when you got on an elevator, pressed the "UP" button and it went DOWN. This is quite upsetting the first time it happens, because your remembered experiences lead you to expect -- "anticipate" -- it to follow the button's label. There are other common experiences of many sorts. (Going up -- or down -- stairs and finding one step, more or less, than was expected.) The point here is that "anticipation" and its related concepts are common occurrences.

What does anticipation consist of? It begins with the existence of a situation where there is a goal to be achieved. The Decision Making Entity examines the memory for ways to reach that goal. There may be an established procedure -- a set of related reference signals -- that needs only be put into operation. Absent such an established method, the DME "looks" for an alternative that appears to result in reaching that goal. It (the DME) selects a promising procedure and uses that remembered set of reference signals. The DME may be using a previously successful procedure, or may be extrapolating from remembered events. Either way, future events are expected, that is, "anticipated."

Oded, same reference, speaks of soldiers being praised for "sticking to the goal." In whatever language or circumstances, this represents an attempt to induce those advised to continue their efforts beyond an ordinary quitting point. If this is done, the ADVISOR gets the credit for any success, and the advisee risks his resources.

END POINT CONTROL.
Powers (930203:3:59 pm EST), Greg Williams (930203) and others.
>End Point control involves comparing what you're experiencing with what
>you want to experience and turning the error into an action that will
>make the error smaller. ... This is how I have thought of configuration
>control, which is truly end-point control.

This statement is followed by discussion of examples presented in terms of "configurations." "Going to Paris" becomes a typical example. However, the reason to go to Paris is omitted. Perhaps a business meeting, perhaps a vacation (to have enjoyable experiences), perhaps for education. These "End-Points" involve much more than "configurations." Anticipation clearly plays a significant part in the decision to "go to Paris." And the DME finds in available memories (including maps, travel agents, etc) the procedures needed. These procedures are then used to provide suitable reference levels as inputs for the systems needed. In this situation, various Skills are needed. Communications to assorted people, handling money and tickets, passports, etc etc. Variables of configuration, sequence and time must be included. And all of these involve suitable control of the lower order muscle skills. Bill, I think this is consistent with your view, but you have stated it in such abstract terms that some of this may be overlooked. It is very helpful to have the concepts of Temporal Variables, Skills, etc available in addition to that of Configurations.

INTENTIONALITY
Chris Malcolm (930204.5:45 pm EST), others
Regarding my discussion of "Anticipation" above, it seems to me that "Intentionality" recognizes that people make decisions (action by something I call the Decision Making Entity) selecting future events/situations to be achieved. e.g., I got in the car with the "intent" of going to the dentist. I "anticipated" little or no traffic and expected the car to perform as it has in the past. I remember the route and the conventions regarding other cars. To me, "Intention" is a Sixth Order Concept -- one uses available Skills to accomplish higher order purposes. Is this a problem?

LANGUAGE -- SOCIAL CONVENTIONS
Bill Powers (930204.1400)
I just want to second your position, particularly:
>The only thing we can say for sure is that each of these understandings
>of the social conventions about language is pragmatically sufficient for
>the task of communication.

and:

>... all we can say about the commonality of these rules and so on is
>that the test of conveyance of meaning and acceptance of forms as being
>reasonably "correct" is not failed.

Language is only one form of communication (very important, of course). There are also vocalizations, gestures, bodily attitude, facial expressions, etc. From a developmental standpoint (of the individual), language is a late arrival.

PHD
Bill Powers (930204.1400)
As a holder of a PhD, sought in part for "practical" reasons, I'd like to report that my life "while getting it" included a minimal percentage of "hell."

MULTIPLE HYPOTHESES
Bill Powers (930204.1400)
>I'm always considering multiple hypotheses where I can think of any. For
>the basic phenomena of control, I can't think of any."

Who is this "I"? From my viewpoint, this is your DME searching your memory for related alternatives. Finding none that you prefer (they don't meet your remembered criteria for acceptability) your DME "plugs in" a combination of remembered skills resulting in your report, above.

ANTI-QUIBBLE

Bill Powers (930204.1500)
I think your "minor quibble" is more serious. The new "aiming point" is the "new target" for the gun crew. The target for the crew is no more no less than that ordered by the commander. To specify it in terms of the preceding target may be a convenient short hand way to communicate the position of the new target.

HIGH LOOP GAIN
It seems to me you are following events around the loop, resembling open loop analysis. Using a time scale including several shots, appropriate to the view of the commander, high loop gain should improve the resulting accuracy. Examining the series of events, we begin with the first shot. It misses by some amount and the location of the impact is reported by the spotter. If the spotter is very sensitive, this location may be reported in feet or inches, although yards might be sufficient. The aim is then adjusted by the crew to whatever accuracy the equipment permits. High gain means that the aim is corrected very precisely. However, the second shot could be off considerably if, for example, there is a gust of wind, the target moved, or whatever. But high loop gain would still tend to minimize the error, instead of creating an over-correction. An over-correction might occur if the gun controls were not properly calibrated. As I understand it, a bracketing procedure is often used to calibrate the gun controls.

Indeed, the "bracketing" concept is useful in any situation (exploration, experimentation) lacking accurate, or reliable, data.

AT THE SAME TIME
Gary Cziko (930206.0225 GMT)
>... no matter how quick the feedback from the spotters, it still seems
>quite different from the continuous control systems that PCT has
>introduced me to. In a "real" continuous control system, perception,
>comparing, and acting are all taking place AT THE SAME TIME.

My statement, "Analysis is influenced by the time scale selected," would have been more clear as, "Whether open loop or closed loop analysis is appropriate depends on the time scale selected."

Closed loop analysis is appropriate for a time scale in which the firing of the gun is completed before the higher order system (the commander's system) can respond. The loop gain has little effect on this analysis because the loop serves as part of the commander's output function. The gain of the loop determines the accuracy with which the output signal follows the reference signal. Loop gain is determined by combining the sensitivity of the spotter with the sensitivity of the gun aiming equipment.

Using open loop analysis, it is indeed "an iterative control process ... like S-R chaining with a reference level," as you observe. However, in the open loop analysis the concept of "a high loop gain" does not apply. There is no "loop" to have a "gain." It particularly does not apply to the gunner alone. The gunner adjusts the aiming equipment according to the correction called for by the spotter. If the report is "100 meters too far," the gunner makes the corresponding correction (perhaps aiming 2 degrees lower); the spotter reports again, etc.

Perhaps this could be called "a qualitative difference between the iterative control and continuous control," but I find it more useful to express it as a "difference in viewpoint." And which view is more useful depends on the purpose of the analysis. The commander's view, with its longer time scale, uses closed loop analysis, the gun-crew-spotters view uses open loop analysis.

MORE "AT THE SAME TIME"
Dag Forssell (930205 11.30)
As you point out, even the electronic signals in wires do not travel with the speed of light. So, how is "at the same time" defined? It is "at a time less than the time of interest." If you care about minutes, the "same time" is seconds. If you care about milli-seconds, the "same time" might be micro-seconds. To jump several steps, if you care about the school year, the "same time" might be within one day's lesson.

This is not a matter of friction or other resistive effects, although those will certainly modify the results.

Quantitative vs Qualitative: Both the open loop and the closed loop descriptions can be described in either qualitative or quantitative terms. The differences are in the viewpoint, and the purpose for the description.

>... As long as the components of the system are dedicated to their task
>and function in a dependable way, they are still control systems."

But why is this suddenly untrue of "social (control) systems?" Perhaps they are subject to an assortment of misunderstandings, mistakes, conflicts, etc, are not these the "disturbances" with respect to which a control system operates? It seems to me that wherever we find such concepts as "corrective action," "quality assurance," "performance evaluation," "achievement tests" and the like, there is at least one control system in action. I am unfamiliar with Deming, though I understand he is regarded as an outstanding authority on business management. It seems to me that business management intrinsically involves an extensive array of inter-related control systems.

Dag Forssell (930205 12.30)
>"CONTINUOUS" vs "AT THE SAME TIME."
>... A behaviorist might listen to Bob Clark and hear him say that this
>is a chain of cause-effect happenings. We in PCT notice the *multiple*
>iterations required to arrive at the target and can see the similarity
>with the iterative calculations of the Marken spreadsheet. ... we see
>the error signals at work, pulling in some (hopefully correct) direction
>and know that the process works well even without perfectly planned and
>executed output functions.

I am not familiar with the Marken spreadsheet, but I can infer the general nature of the demonstration.

The iterations are, of course, steps in the correction process. When observed with a longer time scale, these iterations disappear; at a shorter scale they become more obvious. Purely a matter of viewpoint and choice of time scale for observation.

SEQUENTIAL VS LAGGED CONTROL
Bill Powers (930206.1400 MST)

Your example of the fire hose for "Lagged" control seems to work very well. But I don't think the fire chief cares which form of control it is as long as the water lands where HE specified. The chief uses a time scale of, perhaps, minutes vs the seconds needed for the water to flow.

"Optimizing Control System" using two control systems, each "with its own perceptual signal..." The existence of conflict depends not so much on the nature of the perceptual signal as it does on the relative time scales. Thus the "gun crew plus spotter" is controlling the point of impact of the shell, and so is their commander in assigning the target. If the commander observes excessive spread in the pattern, he may make changes in the lower order system. He might, for example, adjust the position of the spotter to improve his sensitivity. Both systems are concerned with the same perceptual signal, but their output systems operate differently.

As suggested, an "exotic kind of conflict" occurs when the time scales overlap. If the spotter is repeatedly moved to a new position before the operations from the preceding position has been completed, a loss of accuracy (perhaps temporary) results.

Some forms of stuttering provide another illustration. If the individual attempts to correct the formation of his phonemes too soon, ie, before he has completed his word or phrase, stuttering is unavoidable.

Many other examples are readily found.

Regards RKC

Date: Fri Feb 12, 1993 8:36 am PST
Subject: Re: Branching out

[From Rick Marken (930212.0800)] Oded Maler (930211.1100-ET?)

>I think that a paper entitled "An alternative approach
>to motor-control" *can* be written and accepted to BBS. It should
>be written around the lines of the FB too slow discussion.

I like this idea very much. I have only one caveat: it should NOT be written by Bill P., Tom B. myself or anyone else who might conceivably be thought of as a "true believer" (even if they are). I think that it should be written by only one (or two people) for the sake of consistency of tone and style although the ideas in it might be the result of a larger collaboration.

>I think the paper should take advantage of the knowledge of Greg
>and recently Avery of the other approaches to motor control.

Right. I agree. I think the paper can be fairly short (if it is written for PSYCHOLOQUY) -- I would imagine it could be written up (based on existing threads from CSG-L and knowledge of the literature) and submitted to PSYCHOLOQUY by the end of March.

I really think this is an EXCELLENT idea for an experiment. I pre- dict that the paper WILL NOT get accepted if it gives an honest and accurate treatment of the PCT model of "motor control".

I strongly agree with your suggestion that the paper

> need not contain the slogan "B is the C of P" and the pecentage of "internal"
>PCT references should not exceed 10%. It need not speak of a Newtonian
>revolution and all that. All these precautions are needed to convince
>the referees (and some potential co-authors as well :-) that this is
>not a bunch of lunatics.

Absolutely correct. But I think that I pretty much followed all these suggestions in the "Blind men" paper (perhaps not rigorously enough) and it was still rejected.

So I would really like someone else to try getting a PCT based paper published -- one that gives an accurate representation of PCT (whether it uses the terminology or not). I know that it is possible to publish PCT papers that are not really about PCT; this is why Carver and Scheier, Locke, and others manage to publish their stuff -- even though they mention Powers (and even me sometimes) a lot; they manage to present PCT as another approach to explaining "cognitive control" or "guidance BY perceptual feedback". It's no fair to try to publish a PCT paper that is really about "control by constraints" or something currently acceptable to the establishment.

I would LOVE to be proved wrong and see a quality paper on the PCT approach to motor control published in PSYCHOLOQUY (so we could see what kinds of comments it gets and start an honest dialog). So I say "Go for it"! Someone (other than those of us in the PCT lunatic fringe) should write this paper; how about you, Oded? Or Avery? I REALLY want to lose this bet (that it won't be published).

But I STILL think that it can't hurt to advertise the existence of CSG-L to the PSYCHOLQUY readership even if we don't have a paper accepted to "show what we do".

Best Rick

Date: Fri Feb 12, 1993 11:01 am PST
(Avery Andrews 930212.0906)

>On a somewhat different note, I wonder if `feedback' isn't actually a
>dirty word (two four-letter ones). The problem with it is that it
>suggests that the organism's perceptual functions are able to draw
>a distinction between the effects of what the organism does, versus
>the effects of the disturbances, ...

This is a good point. I think it ties in with the thread on "anticipation."

The feedback loop, like my "Structured Control System" construct, exists in the perception of the analyst, not of the living control system being analyzed. One must keep the two points of view separate. If we think of what we analyze to be in a living control system as something that control system perceives and controls, we get into confusion.

Having said that, I'll retract a little with respect to feedback.

The ECSs within a living control system "perceive" only the results of their Perceptual Input Functions (PIFs). The do not perceive their actions, which are simply a bunch of transformations of the ECS's error signal (the difference between the reference signal in the ECS and the perceptual signal in the same ECS.) The ECS doesn't perceive the feedback as such. It perceives the output of the PIF, which, all going well, doesn't deviate much from the reference signal (because, the analyst says, of feedback).

However, we (analysts) do acknowledge that there is an imagination loop, wherein the ECS substitutes the effect of its actions in an imaginary world in place of actions in the real world. It cannot do this without at least implicitly creating the effect that the real feedback loop "should" have--the effects that the actions would have on its perception if the real world behaved like the imaginary one. Disturbances don't enter into this loop, except insofar as they can be independently imagined (i.e. "perceived" as having happened). In the imagination loop, then, the results of actions and of disturbances can be independently perceived.

When it acts in the real world, the ECS does not perceive the results of actions and of disturbances separately. But, and here's the retraction I mentioned earlier, the perception can be compared with the anticipated perception based on the imagination loop without disturbance. In this case, the deviation of the world from its anticipated behaviour could be seen as a perception of disturbance--but note that it is not perception of a real disturbance. The imagination of the anticipated effects of actions on the perception could have been wrong all along, and what is perceived as a disturbance might just be the result of erroneous imagination.

The upshot of this is that feedback can be identified, in a way, by the control system itself, in the coincidence of imagination and reality. It is a process similar to that done by an outside analyst, but is implicit in the ability of the ECS to anticipate, rather than being a separate analytic process.

Feedback can be identified by the outside analyst, but it can only be used by an ECS. And therein may lie some of the confusion Avery identifies.

This discussion would be much cleaner if the place of imagination, planning, and related functions were more precisely identified. But I think that we should not hope for that to happen on the basis of any experimental data for some long time to come.

Martin

Date: Fri Feb 12, 1993 1:38 pm PST
Subject: place cells

Neal,

Yesterday I came up with an idea which, if relatively correct, I think if of great significance (as significant as an idea can be, that is). But before I declare this wonderous thing, I need to check out a fact upon which it is primarily based. (If the fact is wrong, then I'll just have to say "Oh well, if the fact was right, it would have been a great idea.")

OK, so...you have spoken about the existance of cells which have place fields--they are called place cells. The activation of some group of cells X is equivale nt to the Experience of and "object" at place X (recognizing of course that cell s X dont always relate to place X, but in this example they will). It is not th e case, of course, that I simply "see" an object in a certain place and the plac e cells fire accordingly, but rather that the firing of these cells IS the exper ience of the object in that place.

I went through the previous paragraph just to make sure we are both coming from a "subjectivist" approach (which is a misnomer if one is a "subjectivist"...). In other words, I begin with the premise that this world we experience AS real is "really" illusory in the sense that we construct it (again, not really illusory once one accepts a subjectivist approach). That when I experience a box visua lly, it is the same as when I experience a touch on my arm. Its not my arm, per se, which experiences the touch, its the activation of the relevant area of my somatosensory cortex and the place cells presently "covering" that part of my arm that gives me "a feeling right there." Similarly, when I experience visually a box, its the activated relevant visual centers and the activated place cells wh ich give me "a box right there."

OK, now...I would guess that place cells virtually cover all presently relevant areas in "external space." But I see no reason why any cells would activate for the space which is me, taken up by my body. There just wouldn't be any need to since I neither need to know where it is, nor do I need to manuever around it--the two primary reasons why place cells are important. Now perhaps we they do cover the area of our body which we see, say from the chest down. Perhaps. But from the shoulders up, I bet not.

What do you think? Under NORMAL circumstances, do any place cells cover that area taken up by our bodies (our head specifically)? I find when I close my eyes and attempt to imagine and distinguish any two points in space, I can do so. But if I try it for the area where my head is, I cannot. And I do not simply mean imagining a picture of my head--I mean imagining in such a way that I know where it is in 3-D space such that I could direct attention towards it. It does not seem easy and I wouldn't expect it to.

If this is the case (that no place cells cover this area) then what would be the experience of a person if tones were emitted from different regions within a person's brain--could they say where they were coming from at first?

What happens when place fields change to cover this area is the focus of my idea, but I want to check what I've said first before going too far out on a limb. But I will say one thing which isn't a big deal, that when talking, one would go from "I'm talking" (as usual) to "There's a voice coming from inside my head--I suppose that's my voice--but why does it keep talking when I don't feel like I'm telling it to talk."

Speaking of talking, I'm doing alot of it. Let me know what you think cause IO think the next part is really exciting.

Mark

Date: Fri Feb 12, 1993 2:06 pm PST
Subject: comment on place cells

The letter I just posted on place cells was sent to my advisor, but while writing it I thought I could get some interesting comments from the net also, so I sent you'all a copy. I figured that since I would post my conclusion on the matter eventually anyway, I might as well send some of it out now--I apologize that it doesn't directly relate to PCT. But it does indirectly because it centers around the epistemology of PCT--that our "objective" sensory experience is subjective and that the only true objective experience is our own subjective experience. Now there's a quote that will be vague enough to elicit comments--not my point, though. Its a semantic problem, I assure you, not worth clarifying with those who already agree.

Tom, I'm really beginning to see how neuropsych needs to be looking at voluntary phenomena, as opposed to reactive studies, though I could still use some clear examples.

Mark Olson

Date: Fri Feb 12, 1993 2:09 pm PST
Subject: ADAPTIVE MACROPROGRAMMING IN CAD/CAM CONCEPTION

ADAPTIVE #MACROPROGRAMMING IN CAD/CAM CONCEPTION

The works on CAD/CAM systems creation are carried out in the scientific-research laboratory of production engineering optimi- zation of the Georgian Technical University (leader T. Loladze). One of the directions of these works is concerned with making out new technologies and, in particular, conception of adaptive macroprogramming (leader A. Sharmazanashvili), which are the continuation of scientific works started at the department of "Flexible Manufacturing Systems" of the Bauman Moscow State Technical University.

The software were approved and installed in the folowing industries: "TEMP" in Moscow, "ENERGIA" in Moscow, "SPLAV" in Tula, "SPURT" in Zelenograd.

DESCRIPTION

Adaptive macroprogramming is based on the idea of the division of technological operation of machining into structural and parametrical parts, according to the principle of the least affectance of structural elements to the technological disturbances. As a result, structure designing is realized at the stage of technological preparation of production, using the systems of imitative modelling and formation of NC macroprograms and the definition of parameters takes place before the starting of machining process using CNC.

THE CONCEPTION PERMITS:

1. To carry out workable and high reliable NC programs adaptable to the concrete production conditions.
2. To reduce time of NC programs testing at the work shop directly (e.g. at 3 or 4 times). The transition to the complete non-testing technology is possible.
3. To increase productivity and reliability of machining process.
4. To integrate CNC with control systems and make automatic correction of NC programs before the machines are used to cut metal without the operator participation.

THE SOFTWARE intended for the conception realization consist of two parts:

* software for CNC in the form of widened library of unified subroutines

* high productivity integrated environment of technological programming - ISTP/Turbo, realized on IBM PC, intended for working out NC macroprograms and unified subroutines.

ABOUT THE AUTHOR

Mr. Sharmazanashvili Alexander is 32 years old. He is Associate Professor at the Georgian Technical University. During the last years he was engaged in scientific work at the Bauman Moscow State Technical University. He took part in scientific and research works at the enterprises and organizations of military branch in creation of integrated production systems. He is the author of applied software of the CNC of new generation - "Electronika" MC2106 worked out in the designing office "Nauchni centr" in Zelenograd.

In 1991 by decision of the Academic Board of the Bauman Moscow State Technical University, Mr. Sharmazanashvili Alexander was awarded a degree of Candidate of Technical Sciences (Ph.D).

He is a participant of scientific conferences. He is a laureate of the international seminar "Robotica - 90" held in Sozopol Republic of Bulgaria.

He is the author of 15 scientific works.

CONTACT address E-mail: tamazi@tecnex.aod.georgia.su

Date: Sat Feb 13, 1993 5:04 pm PST
Subject: paper

[Avery.Andrews 930214.1130] Oded (930211.1100-ET?) (Rick (930212.0800))

I agree with the main outlines of Oded's proposal. I don't see any reason for a small author list, & am a lot happier with the six-month time-frame than with end of March, since there are quite a lot of issues that I think need to be looked into carefully.

For example, there is an opinion floating around to the effect that removal of feedback leaves behavior more or less intact, while distorting it (by, for example, delay) is extremely disruptive. I suspect that this claim would dissolve under careful examination - some behavior does seem to survive deafferentation with minimal effects (the pointing gestures studied by Bizzi), while other behavior requires a recovery period whose time-scale is consistent, I think I recall, with that required for adaptation to gross sensory changes (such as the prismatic lenses that turn everything upside-down or reverse left and right). My impression is that the feedback- distortion experiments are short-term, not providing enough time for a significant degree of adaptation, so that it might be that there is really no essential difference between the effects. But it will require a substantial amount of bum-on-seat time to check this out adequately.

I see the main theme as trying to make the point that the potential of feedback-based models is being seriously underestimated, so that what it needs is a judicious combination of abstract conceptual discussion (e.g. do something about Kugler & Turvey), and concrete proposals. I think the availability of methods for detecting the presence of various kinds of control systems should be one important theme - e.g. making Marken-style experiments more of a `topic' (as opposed to a `non-topic') than they presently appear to be.

Avery.Andrews@anu.edu.au

Date: Sun Feb 14, 1993 3:57 pm PST
Subject: blind men and learning

I just read Rick's Blind Men paper again (excellent paper, Rick) and I have a few questions (or maybe its one).

How are we defining learning? Big question. I understand everything in Rick's paper just fine, but as I was reading I thought "Oh, OK, just keep everything constant and see what we get when we manipulate D, or Kf, or S*. And while I am at it, can't I say that manipulating Ko (ignore typo) Ko is 'learning.'" I thought this because in a reply of Bill's to me about a week ago on "error control" I understood that the second (reorganizing) system alters the Ko of the ECS. I thought this was learning. But then there is a problem: the phenomena Rick is describing is learning phenomena, isn't it?

When I become "conditioned" to blink to the tone without the puff of air, how did that happen. Rick's paper describes very well how it happens in present time, but how did it get there in the first place? Is Ko altered previous to where Rick's description applies? Or is the tone associated strongly with a puff of air percption such that it is AS IF a puff was perceived, after which Rick's description applies?

Mark

Date: Sun Feb 14, 1993 4:07 pm PST
Subject: learning

I am sending short messages back to back to avoid losing long messages.

Another question about learning: I just read some research in which subjects fixate on a dot and then to another dot to the left (or right) of it. During the saccade, the second dot is moved a consistent amount. This is repeated many times. Eventually, the eye will land on the dot after the saccade even though its presaccadic location is not the same as its postsaccadic location. What happens here?

Is a new reference signal set, representing a different location in space? Or is Ko altered? Either makes sense to me--what's the PCT line on it?

Mark

Date: Mon Feb 15, 1993 6:56 am PST
Subject: Learned Pathways?

I have recently heard on TV that MD's who study pain and deal with its treatment have now said something like it's important to medicate against pain thoroughly enough so that the pathways in the nervous system will not be formed, as they are harder to get rid of once they are there. This goes with what sensible medical practitioners I have dealt with in the past have said: Don't be a martyr, take enough painkiller to keep ahead of the pain and you will in fact need less. However, both my current (Chinese) and ex (Korean) husbands in fact rarely take/took pain medication and the pain didn't/doesn't seem to get to them, or at least nothing was/is shown. (I'm talking about things like headaches, etc. here, not about severe pain in their cases.)

What intrigued me about this is that when I recently saw a network chiropractor several times, she said that present pain is felt in pathways that were created in the past during times of great stress, fear, etc. There seemed to also be some component in their work (though I'm less sure of this) of there being a possibility of their adjustments bringing out the pain or something else for a while but of things then getting better. To some extent this may have been true for me, but I'll spare you the details unless someone wants them. (I saw her for back pain, for the most part; nothing truly horrible, but something generally chronic and not helped by my constantly sitting at a computer :-). But, we all must make our choices.)

If this is of no interest to PCT, just ignore it. But I have a feeling it might be, or it might be so trivial and obvious to the theoreticians out there as to be not worth discussing. However, if it is I'd be interested in hearing how PCT would explain this.

Best, Eileen

Date: Mon Feb 15, 1993 9:20 am PST
Subject: HIGHER LEVELS: I - RKC

[From Bob Clark (930215. 11:55 am EST)] RE: Bill Powers (930201.1900)

(Sorry for the delay in offering the following comments, there were other matters that appeared more urgent.)

In your words:

>I'm taking a different viewpoint: my definitions of levels are meant to
>describe how the world appears from the standpoint of the person
>regardless of the context. When I speak of "system concepts," I'm
>referring not just to things like a self or a personality or a character,
>but to ALL system concepts. To a physicist, for example there exists
>something called physics, a discipline. This is, of course, a perception.
>the entity called physics, I have proposed, is a concept built from a set
>of principles and generalizations, which both provide the material within
>which the entity physics is perceived, and which, as goals, are specified
>by the goals we have for physics -- that is, for what kind of entity we
>want it to be.

>The principles and generalizations, in turn, are built out of a set of
>rational, logical, reasoned mental processes that I call, generically,
>"programs." In a set of programs we can discern general principles; at
>the same time, the principles we wish to maintain in force determine what
>programs we will select to use.

"programs we will select" -- Who, or What, does the selecting? the DME?

Your selection of these higher level structures reflects your extensive knowledge together with the application of a high degree of logical skill and reasoning. However, what about those who are not as knowledgeable? How do they manage? What are the categories, etc that they form and live by? When they interact with other people, what are the concepts they use? How can we talk to them without some common language?

Bill, I am troubled by your move from your Fifth Order, Control of Sequence, to discussion of "concepts." Are these concepts derived from combinations of lower order perceptual variables? If so, how? And which? Does the operation of these concepts include setting reference levels for Fifth Order, and/or lower Order Perceptions? How, and by what is this done?

From the Glossary, BCP, I find:

>"Perception: A perceptual signal (inside a system) that is a continuous
>analogue of a state of affairs outside the system."

Finding no special definition of "concept" in that Glossary, I consult my dictionary.

"concept, n. 1. a general notion or idea; conception. 2. an idea of something formed by mentally combining all its characteristics or particulars; a construct."

I think that's essentially what you mean. What are the perceptual components of "concepts?" It seems to me that this term is too broad and vague a category to be assigned as an Order of Control in the Hierarchy.

Also for "entity," as in "entity called physics" above. Not in the BCP Glossary. Dictionary: "entity, n. 1. something that has a real existence; thing. 2. being or existence, esp. when considered as distinct, independent, or self-contained."

This is how I use "entity" in "Decision Making Entity."

Your view of "physics" seems to differ from mine. To me, a physicist, it is not "a" concept, rather it is a specialized language, including its own special words, syntax, etc. It is an assemblage of definitions, observations, methods, procedures, formulas, derivations, etc etc. I find these in various locations in my memory -- given suitable situations, they are available to select for use, or whatever. In one way or another, any of the lower order perceptual variables may be pertinent. But it does not seem to me to serve as a "concept."

>"What kind of entity we want it to be." From your first paragraph, above.

I don't have any particular "goals" for "physics." It is "set of tools," very useful for certain purposes, but irrelevant for others.

As you recall, my proposal is to assign Control of Temporal Variables to Fourth Order, placing Sequence at Third. Sequences have temporal aspects which are perceivable and controllable. Combinations of Sequences with Temporal Variables, also perceivable and controllable, form Skills. These can provide new sets of perceivable and controllable variables. One Skill can be selected vs another: "Shall we dance the waltz, or the tango?"

Skipping a paragraph in your 930201.1900 post:

>My intention in proposing these levels of perception was to provide a
>framework within which we might understand all human experiences, no
>matter what they are about. If the subject matter is one person's
>experience of other individuals, then what I call "system concepts" would
>correspond to what you term "personality," and perhaps what I call
>"principles" would correspond to your "character," and my "programs" to
>something like "habits" or "abstract skills" or "intelligence."

The "correspondence" you suggest appears to be limited to a similarity in position in the sequence of levels in the hierarchy.

To me, "personality" refers to a group of perceptual variables that have names that are convenient because they are commonly "understood" by ordinary people. They relate to short term interactions and include such perceptual variables as "friendly," "helpful," "dominant," etc. What you call "system concepts" draws pretty much a blank except among those with unusual information and experience. Logical, yes, but the connection with perceptual variables is not clear to me.

To me, "character" refers to another group of perceptual variables. These variables also have names that are "understood" by ordinary people. They relate to identifiable, therefore perceptual, underlying forms of behavior displayed in repeated interactions. Examples include such concepts as "honest," "reliable," "thorough," "careless," -- they are not necessarily favorable. What you call at this point, "principles," in the sense you seem to intend, also draws pretty much a blank except among those with special knowledge as above. Logical, again yes, but what is the nature of the "perceptual variables" from which they are derived, or for which they might provide reference signals?

Similar comments apply to your "programs." "Habits," "abstract skills," "intelligence" I would treat quite differently. To me these are important questions, but are not included in my present comments.

You emphasize: "These (referring to my proposed terminology) are ways of perceiving other people." Yes, but they are also ways of perceiving yourself. We agree that one cannot observe (perceive) ones own acts during the performance of those acts. However this does not prevent their perception by examination of recent (perhaps very recent) memories of those same acts. Near the beginning of this same post, you state:

>To speak of "personality" and "character" is to take an external view of
>someone else's organization. That is, you seem to be looking for levels
>that will apply to "psychological" aspects of a person, to explain the how
>and why of that person's behavior.

This is not at all my own view of what I am doing. I plan to discuss my alternative proposals for higher order systems in some detail in a later post.

Meanwhile, Regards, Bob Clark

Date: Mon Feb 15, 1993 11:29 am PST
Subject: blind men and learning

[From Rick Marken (930215.1000)] Mark Olsen

> I just read Rick's Blind Men paper again (excellent paper, Rick) and I have
>a few questions (or maybe its one).

>How are we defining learning?

>the phenomena Rick is describing is learning phenomena, isn't it?

(I think these questions apply to the "second blind man's" view of behavior, right? Eq. 6 in the paper. Which used operant behavior as an example?).

These are good questions. Discussion of operant behavior is usually found in books on "learning" but the phenomenon I described in my paper (the effects of reinforcment contingencies on response rate) is not a learning phemomenon, though it is the result of learning; the animal must learn (in the sense of developing a control system) how to keep its food (or whatever) input coming in at the desired rate (or as close to it as possible). So the animal learns that it can control the rate at which food arrives by varying the rate at which it presses a bar; My discussion is relevant to what happens AFTER the animal has learned to control its own food input; in that case, the effect on responding of changing the "schedules of reinforcement" -- which, as I mentioned, just changes the feedback function, Kf, from response to controlled variable, is just as predictable and as precise as the effect on responding of disturbance variations in a trackinhg task. The animal "has to" change its responding or the controlled variable will not be controlled. The point of the paper is that these variations in Kf are being interpreted as the effect of reinforcment on responding when they actually work because the animal must vary its responses in order to control the reinforcement.

In PCT, the result of learning is not a new way of responding but a successful way of controlling. In all of the experiments I have done, I don't even start collecting data until the subject has learned to control the variables that I want controlled. In other words, I've never studied learning; just the results of learning (control). In fact, there are very few studies of learning that I know of that come out of PCT. The best I know of is one I've mentioned before; it was done by Dick Robertson et al and published in Perceptual and Motor Skiils or something like that; it tracks a subject's performance as s/he masters a fairly complex control task. What it shows (in the most clear cut cases -- learning, in the sense that it is random reorganization, is not always pretty) is periods of stable control followed by brief periods where performance suddenly deteriorates; after the deterioration, performance is even better than it was before. So the interpretation was that the periods of deteriorated performance are reorganizations where the subject is learning to do the task more effectively; but during this reorganization period the person actually controls more poorly than they had been.

I think learning (reorganization) is wide open for exploration, Mark. It would be nice to have a few basic results under our belt since reorganization is such an important part of everyday adaptation (and, of course, it is probabaly what psychological therapy is all about).

Best Rick

Date: Mon Feb 15, 1993 12:05 pm PST
Subject: Re: blind men and learning; adaptation

[Martin Taylor 930215 14:30] Rick Marken 930215.1000 to Mark Olsen

A nice reply. But I'd like to add a little--

>In PCT, the result of learning is not a new way of responding but a
>successful way of controlling.

I'd like to add "(possibly new)" at the start of the second line.

>In all of the experiments I have done,
>I don't even start collecting data until the subject has learned to control
>the variables that I want controlled. In other words, I've never studied
>learning; just the results of learning (control).

I doubt you achieved that kind of stability. In my incarnations as a psychoacoustician, I found that people continued to learn to detect the signal with no sign of slowdown in their learning rate (measured against the square-root of time) over months of over an hour per day of listening. Many studies in psychoacoustics (and other areas of perception) seem to assume that a day or two, a few thousand (or even a few tens) of trials are enough to bring people to a stable state of ability.

In control studies, improvements in perceptual ability tend to be obscured by the fact that once the gain gets high enough, control is very good. Any effects of learning can't be seen under those conditions. Some control studies are done under difficult conditions (as Tom Bourbon has often pointed out), and then a good model shows a match to the failures of control made by the subject. In such studies, learning should be detectable and modellable.

There are many different places within "classical" PCT where learning can occur. I think I identified twelve, which Bill Powers accepted as valid. Of these, maybe 3 or 4 are likely to be useful in a classical hierarchy. One, output link reconnection, can cause drastic effects when it occurs. Perceptual Input Function refinement can change the detail of a PIF slowly, to match whatever the actions of an ECS are actually affecting. That improves control, but also changes the variable that is being controlled. It does not have drastic effects as it is happening. Bill P says he has studied this form, and finds it useful.

It is unfortunate that the term "reorganization" has been applied indiscriminately to all these forms of learning. I'd prefer to use that word for learning that involves a restructuring of the hierarchy. Learning that involves parameter value changes within a fixed structure might better be called "adaptation" or some such (it's hard to select a good word, because most useful words have connotations in non-PCT approaches, and we probably don't want to evoke those connotations). "Reorganization" should involve a "new way" of controlling, "adaptation" a more skilled use of the same way of controlling.

(Avery.Andrews 930214.1130)

>I think I recall, with that required for adaptation to gross sensory
>changes (such as the prismatic lenses that turn everything upside-down
>or reverse left and right). My impression is that the feedback-
>distortion experiments are short-term, not providing enough time for
>a significant degree of adaptation, so that it might be that there

>is really no essential difference between the effects.

Some of these experiments are by no means short term (as I think of short-term--they could last for days and weeks), and they give very interesting results (such as, with upside-downing prisms, a face seeming right-side up while the cigarette smoke seems to be going downward). J.G. Taylor did some experiments which he claimed to show that the perception compensated for the distortion only to the extent that it was behaviourally controlled. Uncontrolled perceptions remained stubbornly distorted. Without checking, I think much of this is in his 1962 book "The Behavioral Basis of Perception." If not, it is in various papers published around that time.

There is a distinct issue here that keeps cropping up in CSG-L discussions; what one perceives consciously is NOT the same as a controlled variable that is the "perceptual signal" within an ECS. We assume that consciously perceived stuff is built from these perceptual signals, but we can't identify one with the other. To do so leads to confusion. The "perceptions" in an inverting-prism experiment may or may not be a multitude of controlled perceptual signals. They may be a mixture of controlled perceptual signals and uncontrolled transforms of sensory variables. J.G. Taylor's theory would argue that the rectification of the signals depends on the degree to which they are controlled over the duration of the experiment (not all perceptual signals that can be controlled are being controlled at any moment).

Martin

Date: Mon Feb 15, 1993 2:43 pm PST
Subject: trip report, learning, paper

[From Rick Marken (930215.1400)]

First let me give a brief trip report. Bill P. was here and we had a nice visit even though much of the time was spent in a frustrating two day attempt to translate the C version of SIMCON to Turbo Pascal for the Mac. My version of Turbo (1.1) does not seem to agree with System 7. This may have been the reason that we were getting weird "out of memory" compiler errors and syntax errors when there were no syntax errors. So, for the time being, SIMCON (which looks to be quite the PCT modelling tool) is only available for DOS and unix machines.

I did get to show Bill the conflict HyperCard experiments; got some interesting results in the rotation condition (which Bill suggested); basically, Bill solved it cognitively while I came up with some kind of lower order remapping that I was completely unable to describe verbally -- but it worked. Still much more to do -- but I'm going to have to speed up the HyperCard computations considerably if I'm going to impress Mr. Powers.

We also had a wonderful visit with Dag and his wife; kind of a mini CSG meeting in crazy LA. It was all wonderful and perfect except for the rements of a cold I've had for the last several days; damn those intrinsic errors!

Martin Taylor (930215 14:30) comments on my (930215.1000) reply to Mark Olson.

Gee, Martin. I agree with everything you say in that post. Even the stuff about J.G. Taylor; great observations (on adaptation to controlled but not uncontrolled aspects of perception). What's going on here?

Seriously, that was a very wise and worthwhile meditation on learning; we (PCTers) really have to do a LOT MORE work in that area -- both in terms of modelling and research.

Avery.Andrews (930214.1130)

>I agree with the main outlines of Oded's proposal.

>I see the main theme as trying to make the point that the potential
>of feedback-based models is being seriously underestimated

I think this is a wonderful idea; but I still think that we should just go ahead and advertise CSG-L in Psycholoquy and forget about "justifying" ourselves to this audience by submitting a paper with the theme you suggest. I would like it very much if you (or Oded or whomever else) went ahead with such a paper; but I've done this too many times before. I'm sure you guys (and gals too, I hope) can do it better, but I'll bet you almost any amount that you will not be able to get such a paper published in Psycholquy. Even if you write the BEST PAPER EVER on feedback based models and how their value might be seriously underestimated -- you won't get it published if, at any point in the paper, you give an accurate description of the PCT model -- ie. if you correctly describe how a feedback based model actually works.

Even if you are lucky enough to get the paper published (and are figuring out how to spend your winnings) how many of those folks on Psycholoquy will even care about the contributions feedback based models can make to understanding "motor control"? Not many, believe me. So after all that trouble -- and the admitted coup of getting a PCT paper in front of an audience of more than 12 people -- you might get the attention of three people on Psycholoquy, while the rest (who are interested in all the other PCT related things -- cognitive psych, language, inference from text, etc, etc) just say ho hum.

I STRONGLY encourage you (or anyone) to submit a PCT paper to Psycholoquy but, I gotta tell ya, I don't even consider this a bet; I've delt with Harnad for many years and I know for sure that he has no intention of ever letting PCT type stuff into any of his journals. His attitude seems to be that feedback control is history; old hat; and that's that. But PLEASE TRY submitting a paper! I'm definitely through with it, though.

But I would like to have a dialogue with "establishment" psycholgists and I think it might be possible to start it if we could just let it be known that this forum is available. So again, I say "put an ad for CSG-L into Psycholoquy"; so far, it seems like Martin and I are for it; Avery and Oded are against it (until a PCT paper is accepted); Any other opinions?

Best Rick

Date: Mon Feb 15, 1993 7:41 pm PST
Subject: Learned Pathways?

[From Rick Marken (930215.1900)] Eileen Prince says:

>I have recently heard on TV that MD's who study pain and deal with its
>treatment have now said something like it's important to medicate
>against pain thoroughly enough so that the pathways in the nervous
>system will not be formed, as they are harder to get rid of once they are there.

I have heard this too,though not the "pain pathway development" explanation -- which is probably about as good (or crappy) as most other medical theories.

I have also heard that pain control is most effective (and efficient) when it is carried out by the only person who could conceivably control it -- THE PATIENT. Patients who are allowed to regulate their own dosage of pain killer (so that they maintain a perception of NO PAIN or very close to it) experience less pain and use a LOT less pain killer than patients who are given injections by the doctor (at least this is what I heard on "All things considered" -- which is your typical liberal, bleeding heart newscast [perfect for me] -- so i might not be completely correct, but it sounds quite consistent with PCT).

I think a pain control loop like this could be quite unstable (due lags in giving the pain killer and the long time it takes for the pain killer to have an effect on the controlled variable (PAIN) once it is applied. This may be why pain control is so lousy when it is done by an external agent who is also reluctant to give it (I love doctors who refuse to give morphine to terminal cancer patients because the patients might become addicted -- talk about ridiculous -- AND PAIN INFLICTING -- BELIEFS!! But at least they got their standards.). The output (pain killer) is applied in such a way that the controlled variable (pain) starts to oscillate. It would be nice to see (well, maybe "nice" is the wrong word) whether patients who are having pain killer administered by an external agent are always in pain or whether the alternation between high and low levels of pain is just more extreme and higher in frequency. The latter is my control theory prediction -- I'd rather get the evidence from an archive and just let people who are in pain be able to control it.

Best Rick

Date: Mon Feb 15, 1993 8:04 pm PST
Subject: spreading the gospel

i say forget the journal, especially any journal that has accepted feedback models as old hat, or passe' more likely. any missionary knows it is easier to convert the ignorant than heretics just ask the council of nicea or the rejected trio of infidels-bourbon, powers, and marken (and anyone else who have been censored and wound up on psychology's Index. enough of don'ts, now for do's. do continue this digital "grassroots" movement and do establish our own journal. one last don't:do not compromise.

i.n.kurtzer

Date: Mon Feb 15, 1993 11:28 pm PST
From: Jackson

Subject: things with Ray J.

2/15/93 Hello Dag,

We haven't spoken in a while, so I thought I'd drop a quick note. I have the sales tapes to send you...I can see them right here next to the manuscripts I'd like to edit for you. I'm sorry I haven't gotten to either, I'm just way over my head finishing my Master's and trying to keep up with work. Things are going nuts around here, and I'm surrounded by loose ends.

I hope things are going well for you, and I'll get you the tapes soon, with the edits when I can. Sorry for any inconvenience.

Take Care, Ray

Date: Tue Feb 16, 1993 1:20 am PST
Subject: Re: trip report, learning, paper

[From Oded Maler (930216.0925)] Re: Rick Marken on paper

Let me explain my position. I consider the motor control community as serious people knowing far more biology and *relevant* math than myself and than most (if not all) the CSG people. In general their only weakness might be on the mataphysical side, that is, not thinking enough on problems of observability in experiments, not making clear distinction between explicit and imlicit (as-if) computation of the control "output", not understnding the extended meaning of a "program" etc. Sometimes (and I emphasize, *some* times), the knowledge of advanced techniques may prevent from seeing some insight that a less technically sophisticated person might see by naive analysis (this is my own niche in math, btw). I don't think (maybe I'm wrong) that any of these people has a strong *ideological* commitment against feed-back, and if "the paper"is rejcted, it will be for the correct reasons which might be like: "It is nice as an excercise, but it is too naive in the sense that it does not address the real hard problems associated with real dynamics, etc." I'm not saying (because I still don't know) whether such a judgement is correct. Maybe the treatment of disturbances, the simplicity of the computation will compensate for other simplifications.

Just to indicate that feedback is not considered a four-letter word I'll cite from a report of a European research project, "Multisensory Control of Movement" were among the reult they mention: 1) New control principles for visio-motor orienting using visual maps with dynamic remapping properties by velocity or position feed-back. Neurophisiological evidence for the existence of feedback in biological systems. 2) Discovery of the use of extraretinal binocular feedback for percepption of depth. Also there are other established research directions (most notably, "Active Vision") whose models converge toward hierarchical feedback control in this form or another.

About Rick's desire to make a "come back" to the psychological community in the large, I must admit that I don't share this passion. The trade-off between quality and quantity of audience is not attractive - why not try to get into the popular press? I also don't see what you can sell these people except for telling them that their discipline is ill-founded? The higer-levels are speculations which should be discussed among philospohers rather than among people who pretend to do experimental science. Except for confusing some already confused students, I don't see what you will get from it (of course I forgot your prevertic satisfaction from rejection :-) I really understand your motivation - I also sometimes submit papers to mainstream theoretical CS conferences with high rejection probability, as a way of showing them "my problems are more interesting than yours".

But Psycholoqy accept ads, it's ok if you try to hunt some lost souls from their readerships.

--Oded

p.s. I enclose a message I found at comp.ai

From: bright@ee.ualberta.ca (Dan Bright)
Subject: Help: Robot Architectures
Organization: University of Alberta Electrical Engineering
Date: Thu, 11 Feb 1993 07:19:51 GMT

Hello,

I am presently working on a design approach for the lowest level of a self-organizing, hierarchical, robot control architecture. I believe that the architecture within which I am working is a variant of Brooks' subsumption architecture. However, I need to clarify this, and find out who is working with similar concepts.

The architecture I am using is based on a hierarchical communications model, where communications is restricted to adjacent layers. Within the hierarchy, each layer is viewed as an autonomous machine which:

1. operates in an environment provided to it by the next lower
   level of the hierarchy.

2. provides the layer above it with a more abstract environment within which to operate.

```
             ^      |             Abstract Environment (Symbolic)
      sensor |      |  actuator   --------------------
             |      V
          -----------                     ^
          |         |                     |
          |         |                     |
          -----------                     |
             ^      |                     |
      sensor |      |  actuator           |
             |      V                     |
          -----------                     |
          |         |                     |
          |         |                     |
          -----------                     |
             ^      |                     |
      sensor |      |  actuator
             |      V            ---------------------
                                 Physical Environment
```

Any relevant comments or references would be greatly appreciated.

Thanks, Dan Bright

Date: Tue Feb 16, 1993 1:41 am PST
Subject: paper targetting

[Avery.Andrews 930216.2030]

It may be the BBS/Psycholoquy is not the best target - Journal of Motor Control might be more suitable. For one thing, that's where the perceived problem that we might address is (motor programming vs. action theory). Also I'm not sure that the BBS peer commentary is exactly what PCT needs at this point - the time and space constraints make it hard for the commentators to do much besides shoot from whatever positions they happen to have pre-prepared. And contra recent Oded, the stuff that appears there does not strike me as cosmically more sophisticated than what Bill or Greg can handle.

Avery.Andrews@anu.edu.au

Date: Tue Feb 16, 1993 1:56 am PST
Subject: Re: paper targetting

[From Oded Maler 930216.1040-ET]
*
* [Avery.Andrews 930216.2030]
*
* It may be the BBS/Psycholoquy is not the best target - Journal of Motor
* Control might be more suitable. For one thing, that's where the
* perceived problem that we might address is (motor programming vs.
* action theory). Also I'm not sure that the BBS peer commentary is
* exactly what PCT needs at this point - the time and space constraints
* make it hard for the commentators to do much besides shoot from whatever
* positions they happen to have pre-prepared. And contra recent Oded, the
* stuff that appears there does not strike me as cosmically more
* sophisticated than what Bill or Greg can handle.
*

I guess you are right, I hope also concerning the relative sophistication - I am talking *before* making the investment of looking into the deatails, just about surface impressions.

Btw, have you cam across the survey of T. Flash, The organization of human arm trajectory control, 1990 ?

--Oded

Date: Tue Feb 16, 1993 3:30 am PST
Subject: Re: paper targetting

[Avery Andrews 930216.2230] (Oded Maler 930216.1030]

My impressions aren't all that deep either, and I suspect that there's been recently a big upsurge in level of sophistication of the mathematical lines people are talking, but it's hard to tell how much better thought out the actual concepts are. Kugler & Turvey certainly sound real impressive when they talk about rhythm, but it's hard to see how they could have anything useful to say about how people manage to make & eat their breakfasts.

One indication for trying JMB is that they do occasionally publish articles that are more philosophical than factual, such as one from 1986 or thereabouts, in which the author explained why he thought that theoretical positions in his area were use. So a politely worded article to the effect that closed-loop theories might not be quite as terminally dead as they are sometimes said to be might have a chance.

I haven't seen the T. Flash article - where is it?

Avery.Andrews@anu.edu.au

Date: Tue Feb 16, 1993 4:58 am PST
Subject: Re: paper targetting

From: Oded Maler 930216.1330] Avery Andrews 930216.2230

* I haven't seen the T. Flash article - where is it?

It's chapter 18 in J.M. Winters and S.L-Y. Woo (Eds.), Multiple Muscle Systems: Biomechanics and Movements Organization, Springer, 1990. I don't have the whole volume, but I guess it contains a lot of others relevant surveys.

The article contains, in particular, description of experiment in arm trajectory modification, that is, the subject is told to point to some target, and while moving (not seeing his own hand to prevent visual feed-back) he is told to point to another target. One suggested explanation (using terminology of "plans", but that's not the point) that the resutling trajectory is not a result of "aborting" the first one and starting the new one, but that it is a result of "vector addition" of the initial trajectory with a time-shifted trajectory from the origin to the new target. It might be interesting to see how it such things are explained based on Little Man's phisiology, and how the measured velocity profiles comapre to simulation LM's simulation results.

--Oded

Date: Tue Feb 16, 1993 8:10 am PST
Subject: Re: paper targetting

[Martin Taylor 930216 10:45] (Oded Maler 930216.1330)

>The article [chapter 18 in J.M. Winters and S.L-Y. Woo (Eds.), Multiple
>Muscle Systems: Biomechanics and Movements Organization, Springer, 1990.]
>contains, in particular, description of experiment in
>arm trajectory modification, that is, the subject is told to point
>to some target, and while moving (not seeing his own hand to prevent
>visual feed-back) he is told to point to another target. One suggested

>explanation (using terminology of "plans", but that's not the point)
>that the resutling trajectory is not a result of "aborting" the first
>one and starting the new one, but that it is a result of "vector
>addition" of the initial trajectory with a time-shifted trajectory
>from the origin to the new target.

Sounds remarkably like the cover article in the June 19, 1992 issue of Science (Georgopoulos et al., The Motor Cortex and the Coding of Force, v 256 p 1692-1695) except that they used monkeys and looked at the firings of motor cortical cells. The firings were related to what in PCT would be considered the error in the perception (it was a force-to-cursor control rather than a position-to-cursor relationship). The authors didn't look at it from that point of view, but the results look very nice in PCT terms.

When I mentioned this paper at the time, Bill P responded that he wasn't particularly excited about it. The interest would have been if the nerophysiological results had not agreed with the PCT position. But it seems to me to be just one more line of support that can be used in any propaganda article, wherever it might be directed.

Martin

Date: Tue Feb 16, 1993 9:10 am PST
Subject: Re: paper targetting

[From Rick Marken (930216.0800)]

So I take it that I can add Isaac K. to my side (advertising CSG-L in Psycholoquy -- no paper submission first)?

Oded Maler (930216.0925)

>Let me explain my position. I consider the motor control community
>as serious people knowing far more biology and *relevant* math
>than myself and than most (if not all) the CSG people.

This is all probabaly true; what they don't seem to know is anything about the nature of control, how it works, what it implies about how we must study behavior and why it MUST BE the basic organizing principle of living systems.

>Just to indicate that feedback is not considered a four-letter word
>I'll cite from a report of a European research project

I have no doubt that "feedback" has its place in "conventional" motor control models. It is just not recognized as fundemental. PCT shows that it is; all behavior is organized around the control of perceptual variables (feedback).

>About Rick's desire to make a "come back" to the psychological
>community in the large, I must admit that I don't share this
>passion. The trade-off between quality and quantity of audience
>is not attractive - why not try to get into the popular press?
>I enclose a message I found at comp.ai

I'm not interested in a "come back" (I was never "there" anyway). I just want a dialogue with people who want to understand the same phenomenon that PCT wants to understand: human behavior -- from finger movements to political movements. We (on CSG-L) often talk about how "conventional psychologists" deal with such and such phenomenon -- and then we say why it's wrong. Why don't we invite these conventional psychologists to participate in the dialogue; not to brow beat them (hopefully) but to learn what they think and why they think it. Then we might try to explain why PCT provides a better way of approaching their problems (if it does). That's why I want to advertise CSG-L in Psycholoquy. I think the quality of the dialogue on CSG-L might actually improve if we get some bright representatives of the "opposition" to participate.

>I enclose a message I found at comp.ai

>1. operates in an environment provided to it by the next lower
>   level of the hierarchy.

>2. provides the layer above it with a more
>   abstract environment within which to operate.

Wow. Great find! This fellow should definitely be on CSG-L. I'll send him a note later today -- would you do it too in case I forget?

Avery.Andrews (930216.2030)

>It may be the BBS/Psycholoquy is not the best target - Journal of Motor
>Control might be more suitable.

We (Bill Powers and I) tried JMB (the "Levels of intention" paper -- reprinted in "Mind Readings" was rejected by JMB after a LOT of interaction with some of THE most ridiculous reviewers of all time.) But maybe you can do it better; I betcha that you won't have any better luck getting a PCT based paper into JMB than into Psycholoquy. Again, I very much hope that I am wrong and that my PCT publishing failures are the result of my own incompetence.

Best Rick

Date: Tue Feb 16, 1993 11:23 am PST
Subject: Paper; Psycholoquy

[from Gary Cziko (930216.1735 GMT]

Avery Andrews et al.:

Concerning the "motor control" paper, it may be a good idea to say something about the vestibular-ocular-reflex which seems to run in "real time" as an open-loop system (although reorganizes (sorry, Martin, adapts--I do like the distinction) in order to keep retinal images stable). This would show that you (we, whoever) are not opposed in principel to open-loop systems where in fact they seem to work (in the protected confines of our nearly disturbance-free eye sockets). But this may in fact be the only open-loop behavioral system in humans (Wayne Hershberger and his students should be able to help with this).

Rick Marken (930216.0800) says:

>So I take it that I can add Isaac K. to my side (advertising
>CSG-L in Psycholoquy -- no paper submission first)?

You add me, too.

Why don't we use the resources of the net to compose our "ad?" If this is refused by Psycholoquy, there is nothing that can stop any one of us from sending it anyway to the 300-or so "public" subscribers to Psycholoquy.

Rick, could you send to the net examples of such ads that Psycholoquy has run?

>Wow. Great find! This fellow should definitely be on CSG-L. I'll
>send him a note later today -- would you do it too in case I forget?

I will also send him the intro material on CSGnet.--Gary

Date: Tue Feb 16, 1993 11:44 am PST
Subject: Re: re: re: control article

[Hans Blom, 930216]

How pointing to a possibly interesting reference can lead to an avalanche of opinions! Here's some more...

(Martin Taylor 930211 10:00)

>One of Bill's points repeated over and over is that there is NO requirement
>for, or consideration of, linearity in the control loop--ANYWHERE.
>There is a problem with non-monotonic relationships, which has been
>addressed in several exchanges, but monotonicity is the strongest
>requirement put on what one might call "control in the normal Bill Powers
>sense."

I did not mean to imply that a linear relationship is required in linear control schemes, only that the non-linearity can be neglected. Let me be more specific. If the plant to be controlled has a non-linear but monotonic input-output transfer function, it can be controlled by a linear control- ler. It is just as if the system's loop gain changes depending upon the point of operation. If the system has a low gain, the controller needs a high one and the other way around. A design that is based on the highest gain that the system has will work, but it will be sluggish in those regions where the system's gain is low. Sometimes that won't matter, but in high-performance systems it DOES matter. Then it is necessary to know (a priori or through learning/adaptation) and account for the nature of the non-linarity in the control process. That is where non-linear control schemes come in. In cases where the non-linearity is non-monotonic, such as a hysteresis, linear control schemes break down. My point was, however, that as soon as ANY type of non-linearity exists, only non-linear control schemes will lead to peak performance.

> We have indeed talked about ECSs with dead zones, in which the error is
>effectively zero over some range of (perception-reference) (this can be
>done in the comparator or in the output gain function). Your car not going
>over the edge can readily be handled this way, although a function more

>like a cubic would seem more useful than a threshold function.

This I do not understand. Are you actually suggesting some type of nonlinear control scheme?

>The fundamental point of PCT is "if you can't perceive it, you can't do
>anything about it." Everything you do is related to something you can
>perceive. Everything you do may have side-effects that you can't or don't
>perceive, but that is of interest only to someone else, not to you, even
>though it might kill you in the end.

The fundamental point of control theory is "if you can't predict it, you can't do anything about it". All control schemes depend, implicitly or explicitly, on a model of the system to be controlled (where model is used in the meaning of a description of the relevant aspects of the system). The model, whatever form it takes, formulates the output consequences of the inputs. You can do three things with a model:

```
      input     -------------   output
      ----------|   model   |---------
                -------------
```

1. given input and model, we can calculate the (most likely) outputs; this is called prediction;
2. given model and output, we can calculate the (best) inputs; this is called control;
3. given input and output, we can calculate the (best) model; this is called system identification.

>And yes, references for avoidance are easily accomodated in the theory.
>A saturating nonlinearity in an ECS with a positive feedback loop will do it.

This is unclear to me. Please define 'avoidance' in terms of the model.

(Rick Marken (930211.0800))

>I will try to answer your points but first I must ask if you have ever read
>Powers' books ("Behavior: The control of perception" and "Living control
>systems I & II")?

I have read BCP and the Byte articles. I could not find the others in our library. Does that qualify me?

>Have you seen the ARM demo? Plenty of non-linearities in that "plant"
>(the muscles, environment, etc). Have you seen Bill's 1978 Psych Review
>paper (reprinted in LCS I) where (among other gorgeous things) he shows
>how people (and the PCT model) can handle a CUBIC "plant" function (about
>as non-linear as you can get) when controlling a cursor (the function
>relates handle "outputs" to cursor "inputs".

PEOPLE are very good (but often highly non-linear) controllers. Moreover, it is my perception that people have a whole range of control schemes and frequently even apply the appropriate one at the appropriate time. This is a continual source of amazement (and envy) for control engineers who generally do much worse.

>   organisms control PERCEPTION, not output.

I know by now what you mean by the mantra 'organisms control perception'. As so often with jargon, it is an abbreviation for a whole philosophy and only understandable for those who have gotten to know that philosophy. It is right, from a certain perspective. From another perspective, organisms control their outputs. I find it hard, in a control loop, to see one apart from the other. But, of course, sometimes you concentrate on the one, sometimes on the other. Very often, the output is controlled as well, for example in cases where different actions are possible (steak or salmon?), all leading to similar perceptions (great food!). Then you actively have to choose between outputs ("I would like ...").

> The non-linearities are one reason why anything other than
>a control organization never evolved.

What is the control organization of a virus over and beyond that of a hydrogen molecule?

>>Perception is not the only human capability that we depend on to control
>>our behavior.

>Perception does not control our behavior; perception is CONTROLLED by
>actions; controlled perception IS behavior.

Perception is controlled by actions; actions are controlled by perception. Remember the loop!

>>Sometimes memory will do:

>Please, show me the diagram of this model; the only sense I can make of
>it is that references can be set based on remembered perceptions.

See the diagram above, where I referred to system identification. The assumption is that humans, just like adaptive control systems, have some type of 'correlator' or 'associator' built in that tells how two perceptions relate. Once something is learnt, it can be used to predict ('what happens if I do X') or to control ('what should I do to achieve Y'). In this view, learning or 'reorganization' is not a random process, but a process much like curve-fitting in statistics. Learning is necessarily what Skinner calls 'superstitious'; it has nothing to do with truth, only with correlations. The difference between superstition and 'true' belief is, of course, a qualification from an external point of view. One person's belief is another person's superstition.

>>a child will stay away from a hot stove after having been 'bit' by it only
>>once.

>How does it "stay away"; my guess is that after the experience of
>being burned the child starts controlling it's distance (a PERCEPTION)
>from the stove differently -- it set's it's reference for distance
>at a value that might be called "far" rather than "close".

This is what I contest. There is no reference for distance, no single 'far'. Control systems work with _numerical_ values. Therefore 'at least as far as X' is required, with the value of X depending on other perceptions.

>> I do not dispute that we have reference levels and
>>that we use our perceptions to get us close to them. I just want to add
>>something like 'negative reference levels', things to stay away from.

>We don't "use our perceptions" to get close to reference levels; we use
>the effects we produce on the environment (linear or not) to influence
>our perceptions to keep those perceptions tracking internally
>specified reference levels. 'Negative reference levels' are just
>settings for perceptions that we give names to like "avoidance".

My reference for pain is zero. Having a zero perception of pain, however, does not tell me how far away I am from pain. This is a serious matter in drug delivery systems that infuse a pain killer ('pain' can be established from a number of easily measured signals; the stability of the blood pressure and the heart rate are examples). Unsophisticated systems occasionally deliver huge overdoses. Regrettably, only unsophisticated systems exist for unconscious patients. If the patient is in control, drug delivery proceeds beautifully.

>>As a control systems designer, I must seriously object. We do not create
>>control systems 'in the hope that' they function correctly; hope has no
>>place in the model. We do not rely on things going right only 'usually'.

>You can produce a control system that controls, say, it's internal
>termperature perfectly, compensating for normal and even extreme
>disturbances perfectly. Then the system is hit by a meteorite; no
>more control.

I meant something different. A control system is normally designed from specifications. Specifications state which situations the system is to (be able to) handle and therefore also which not. Completeness is impossible, because it is impossible to imagine EVERYTHING that might occur, and also because it would be too costly. If meteorites are a cause for concern, a meteorite-resistant control system should be built.

>>What is so difficult in accepting that there are some things that we want
>>and other things that we do not want? If a model cannot handle negation,
>>too bad for the model.

>Isn't "not wanting" just a way of saying that you want something at
>a particular level -- like zero (not wanting to taste lima beans -- ie.
>wanting zero amount of that taste perception) or 1000 (like not wanting
>to be near a lawyer -- ie. wanting to be 1000 miles away from the nearest
>lawyer).

Are you serious? Wanting to be exactly 1000 miles away from the nearest lawyer is a goal not easily achieved!

(Avery.Andrews, 11-FEB-1993 23:22:22.23)

>Hans, Since you seem to be an actual control engineer, I have a question that
>you might be able to answer. In my readings in the psychological

>literature, I've occasionally encountered claims to the effect that
>feedback is control `technically' restricted to situations where what
>is controlled is the (in some sense which I don't really understand)
>`direct' output of the plant. For example, regulation of muscle-length
>by the spindle organs would qualify as control, but control of the
>distance between the lips by moving one a greater distance when the
>other is disturbed would not. My question is whether this is an
>actual doctrine recognizable from your technical education, or just
>an idea that certain people have picked up somehow.

This is from Kuo's popular textbook "Automatic Control Systems" (page 2):

> Regardless of what type of control system we have, the basic ingredients of the system can be described by
>
> 1. Objectives of control.
> 2. Control system components.
> 3. Results.
>
> ...
>
> In general, the objective of the control system is to control the outputs in some prescribed manner by the inputs through the elements of the control system. The inputs of the system are also called the actuating signals, and the outputs are known as the controlled variables.

The difference you're looking for may be this: some of the control system's results (outputs) are prescribed (the oil refinery's gasoline production volume), some are not (the volume of its exhaust fumes). A control system is designed to achieve the former; the latter are incidental by-products.

(Avery Andrews 920212.1024)

>I wonder of some of the friction between control system engineering and
>PCT might be due to the fundamentally different aims of the two subjects.
>A control system engineer presumably needs to be sure that their design
>will work properly, before advising anybody to spend money to build it.
>It is therefore a high priority to find a convenient subset of the
>possible control systems that have nice mathematical properties that
>allow things to be proved about them.

There is no friction here, in my opinion. Kuo [ibid, page 12]: "For linear systems there exists a wealth of analytical and graphical techniques for design and analysis purposes. However, nonlinear systems are very difficult to treat mathematically, and there are no general methods that may be used to solve a wide class of nonlinear systems." If possible at all, we design linear systems. But sometimes, when the highest performance is required, in my case in medical applications, those just aren't good enough. For a control engineer, that is most regrettable. For patients, it is not. And we might be one ourselves, one day...

>PCT on the other hand is concerned with identifying the control systems
>that natural selection and learning methods have come up with, and these
>operate completely differently from the way engineers do - they don't
>care at all about proofs and nice mathematics, but simply experiment
>massively in all directions, and cull the variants that don't work.

Right! A most fascinating research area, that I'm studying right now. What facinates me most is how _efficient_ the ones who succeed are, what clever tricks they use to do the things they do in the fastest and/or most energy-efficient manner. Great little engineers!

Best, Hans

Date: Tue Feb 16, 1993 1:09 pm PST
Subject: Mantra Schmantra

[From Rick Marken (930216.1200)] Gary Cziko (930216.1735 GMT) --

>Rick, could you send to the net examples of such ads that Psycholoquy has run?

I managed to throw away a rather long Psycholoquy ad for a new electronic journal called PSYCHE. This journal is refereed and it is about research on consciousness but otherwise it seems very much like CSG-L. So Psycholoquy DOES take ads for new groups of interest; I think we should do it. Now that I think of it, I bet Harnad would not even let an ad for CSG-L into Psycholoquy. So maybe the CSG-L ad idea is moot anyway. I'm tellin' ya; this PCT stuff is subversive; where's Joe McCarthy when you need him?

Hans Blom (930216) --

>I know by now what you mean by the mantra 'organisms control perception'.
>As so often with jargon, it is an abbreviation for a whole philosophy and
>only understandable for those who have gotten to know that philosophy. It
>is right, from a certain perspective. From another perspective, organisms
>control their outputs.

I think this is the crux; the point where control engineers (and just about everybody else) seems to miss the boat. "Control of perception" is neither a mantra nore a philosophy -- it is a fact about control system operation; one that a control engineer can usually ignore because s/he is most interested in the objective correlate of the perceptual variable -- the controlled variable in the environment. Psychologists can't ignore this fact because they are not designing the system; they are trying to understand it and objectivized side effects of the control process, though visually interesting, are irrelevant to understanding behavior.

I'm just too tired to try to prove (for the um-teenth time) that the perceptual signal is the controlled variable in a negative feedback control loop. I can only ask "From what perspective do organisms control their outputs"? The only perspective I can imagine is the one taken by the control engineer -- where the word "output" might refer to the objective correlate of the perceptual controlled variable itself (for example, the air temperature near the sensor of the thermostat might be called the "output" of the thermostat since this variable depends (in part) on the heat generated by the thermostat's furnace. This variable could be considered "controlled" but it is not really; the only variable being controlled by the thermostat -- EVER -- is its perceptual variable, in the form of an electrical voltage or current, usually).

The fact that it is the percetual variable that is controlled in a negative feedback look is not a mantra, a philosophy or a political stance. It is a fact of control system operation; it is an observation about control systems that could have been made in the 1930s. It was made in the 1950s (originally) by William T. Powers. It is an observation of such revolutionary significance to the life sciences that none of those sciences want to hear about it because it would put them out of "business as usual" (it will have virtually NO effect on control engineering, by the way). It is the observation that is the very essense of CSG-L -- it's why were here.

If you think that this observation is incorrect; if you have a way of showing that it is not always the functional equivalent of the perceptual signal in a control loop that is controlled -- that sometimes, for example, it is the functional equivalent of the output variable that is controlled -- then we would REALLY like to know about it; I would like to know about it because I have spent a good deal of my life trying to show that behavior is the control of perception; if it's sometimes the control of output or error or disturbance or some other variable in the loop the I'm going to have some 'splainin' to do.

Best Rick

Date: Tue Feb 16, 1993 3:08 pm PST
Subject: Re: control article

[Martin Taylor 930216 16:20] Hans Blom, 930216

How your posting brings back memories! It could have been one of mine, 18 months ago or so. I don't know whether I've learned new truths, or just adapted to the CSG-L milieu. Anyway, I'll try to answer as best I can from the way I now understand PCT. Since you ask.

>...If the plant to be controlled has a non-linear but monotonic
>input-output transfer function, it can be controlled by a linear control-
>ler. It is just as if the system's loop gain changes depending upon the
>point of operation. If the system has a low gain, the controller needs a
>high one and the other way around. A design that is based on the highest
>gain that the system has will work, but it will be sluggish in those
>regions where the system's gain is low. Sometimes that won't matter, but in
>high-performance systems it DOES matter. Then it is necessary to know (a
>priori or through learning/adaptation) and account for the nature of the
>non-linarity in the control process.

The only problem I see in this is the word "know" in the second-last line. There are two senses in which this "knowledge" can be used within PCT (as I understand it). One is in the pattern of connections (including weights, if applicable) among ECSs and in the form of the perceptual input functions of ECSs. The other is in the output function itself (including any distribution of the output to other ECSs, if that is something that the output function can affect). The first form relates to a connectionist distributed knowledge that cannot be overt, even in the highest levels of a control hierarchy. The second form could possibly be overt, and may be overt at high levels. Only the second form, when overt, corresponds to the everyday sense of "know." Usually, "is able to do" is a reasonable substitute for "know how to do," in cases where the distinction matters.

>In cases where the non-linearity is non-monotonic, such as

>a hysteresis, linear control schemes break down.

This, I think, is necessarily the case at the category level and above. You cannot have categories without hysteresis. At least that's my belief. And if this is so, hysteresis is at the foundation of symbolic control.

>My point was, however, that as soon as ANY type of non-linearity exists,
>only non-linear control schemes will lead to peak performance.

So? Where's the issue with/against PCT?

>> We have indeed talked about ECSs with dead zones, in which the error is
>>effectively zero over some range of (perception-reference) (this can be
>>done in the comparator or in the output gain function). Your car not going
>>over the edge can readily be handled this way, although a function more
>>like a cubic would seem more useful than a threshold function.
>
>This I do not understand. Are you actually suggesting some type of non-
>linear control scheme?

I thought so, but given your initial comment:
>I did not mean to imply that a linear relationship is required in linear
>control schemes, only that the non-linearity can be neglected.
I'm not sure whether it would be non-linear to you.

What I mean is that for small differences between reference and percept, the output is effectively zero, but if the distance gets large enough, the restoring "force" (i.e. output) gets larger according to a cubic law. Dynamic gain increases as the square of the error, if the function is cubic. But the output could instead get larger according to a step function, a positive step at some positive value of error and a negative step at some negative value of error. Or it could be zero up to a threshold error and then increase linearly. There are lots of possibilities. The situation is one in which a small deviation from the centre of the road is no problem, but a large one is disaster, so increasing gain with increasing error is a reasonable kind of control law. I don't know whether you call this a type of non-linear control scheme. I do.

>The fundamental point of control theory is "if you can't predict it, you
>can't do anything about it".

If you can predict exactly what the results of actions will be, you don't need control. If you can't predict at all how your actions will affect the world, you can't perform control. If pushing and pulling are both equally likely to bring the object closer or further, and it is going to move the same way even if you do nothing, you can't bring the distance to where you want it. Control is necessary and possible when the effects of actions have some effect on the world that results in perceptions that are more predictable than would be the case if you didn't act. If you can't perceive what the state is (allowing for the inclusion of its current derivatives, if they exist, and whatever predictions your model of the world can provide), then control is not possible. I reiterate:

>>The fundamental point of PCT is "if you can't perceive it, you can't do
>>anything about it."

>The model, whatever form it takes, formulates the output consequences
>of the inputs. You can do three things with a model:

```
>
>      input     -------------   output
>      ----------|   model   |---------
>                -------------
>
```

>1. given input and model, we can calculate the (most likely) outputs; this
>   is called prediction;
>2. given model and output, we can calculate the (best) inputs; this is
>   called control;
>3. given input and output, we can calculate the (best) model; this is
>   called system identification.

I take it that the line you have marked "input" is the actions you impose on the world, and "output" is what goes into your sensory apparatus? If so, there isn't a great deal of problem here, though we might quibble on some points. The model, for example, will incorporate several sources of uncertainty. You could have added an input to the model marked "disturbance". The word "(best)" might be a problem, but not one I'd like to pursue here. But for the most part, OK.

>>And yes, references for avoidance are easily accomodated in the theory.
>>A saturating nonlinearity in an ECS with a positive feedback loop will do it.
>
>This is unclear to me. Please define 'avoidance' in terms of the model.

Avoidance is that the ECS brings its percept to a value different from the reference in either direction. I assume a set of connections that provides positive feedback, and an output gain high near a zero error and dropping to zero when the absolute error is large.

In answering Rick (930211.0800), you get into questions of decision. There are at least three aspects to decision. And again it relates to "overt" and "covert" operation.

(1) If two actions are incompatible in the world, the control loops that involve them are mutually inhibitory. Only one of the relevant perceptions (or neither) will actually be controlled. We have had some discussion of this in the past. I like to look on it as a question of loop impedance, in which the coupled real-world systems affect each other's impedance and hence the gain of the loops in which they participate.

(2) The second aspect of decision is in the distribution of output from an ECS. It is conceptually possible for an output function to alter the distribution of its outputs as a function of the error. That could be seen as decision (I am far in space and near in time to somewhere I want to perceive myself, so I take the car, whereas if I were near in space and far in time, I would walk).

(3) A third aspect of decision is in the program level. As I see it, a program is a set of bifurcations (or splits of higher multiplicity). Each bifurcation selects a reference sequence for lower ECSs. That, too, is decision.

>I know by now what you mean by the mantra 'organisms control perception'.

>As so often with jargon, it is an abbreviation for a whole philosophy and
>only understandable for those who have gotten to know that philosophy. It
>is right, from a certain perspective. From another perspective, organisms
>control their outputs. I find it hard, in a control loop, to see one apart
>from the other. But, of course, sometimes you concentrate on the one,
>sometimes on the other. Very often, the output is controlled as well, for
>example in cases where different actions are possible (steak or salmon?),
>all leading to similar perceptions (great food!). Then you actively have to
>choose between outputs ("I would like ...").

OK. This is the hard part, as I know from experience. As an analyst, you are quite right to say that it is hard to see one part of the loop separately from another. The information about the loop behaviour is available all around the loop. But from the viewpoint of the organism that uses the loop, things are quite different. The perceptual signal is your only way to determine what is going on in the world. You can't detect your outputs, and you can't control them. What you can control is the perceptions that result from your outputs, and that inherently affects the outputs themselves.

To the analyst it looks as if the outputs are being controlled, but that is because the analyst can perceive those outputs and not the organism's sensory inputs in his or her own control hierarchy. The organism tensing a muscle can perceive the inputs from sensors within the muscles; an organism moving an arm can perceive joint angles from various kinaesthetic sensors (I think). But the outputs themselves are not sensed *as outputs* and therefore cannot be controlled.

It's not a philosophical position that "organisms control perception, not output." It's a simple statement of fact, about which not much can be done except to note its necessity.

When you talk about choosing between outputs (salmon or steak), you are talking about a whole hierarchy of controlled perceptions involving learned connections that get you to perceive one or the other set of senstions. The connections ARE on the output side--without that, you could do nothing--but they are not controlled; they are just THERE. How they got there is a separate issue (reorganization in the PCT context).

>Perception is controlled by actions; actions are controlled by perception.
>Remember the loop!

Perception is controlled as a consequence of the loop. Neither of the statements in the first line is true. You have to change both "controlled by" into "affected by."

>See the diagram above, where I referred to system identification. The
>assumption is that humans, just like adaptive control systems, have some
>type of 'correlator' or 'associator' built in that tells how two percept-
>ions relate.

I'm sympathetic to this. It makes for efficient control. See BCP pp224ff.

>Once something is learnt, it can be used to predict ('what
>happens if I do X') or to control ('what should I do to achieve Y'). In
>this view, learning or 'reorganization' is not a random process, but a

>process much like curve-fitting in statistics.

There are different kinds of learning within the PCT framework. Some are necessarily random, some not. For instance, the connections of the sensory-perceptual functions are identically those of a multilayer perceptron in the simplest case (Perceptual Input Function is a weighted sum followed by a squashing function like a logistic). Any learning algorithm applicable to a MLP could be applied, the trainer being the desired (reference) set of perceptual signals. That IS very like curve- fitting in statistics (I think John Bridle argues that it is mathematically identical).

But it is not possible to curve fit when the question is about what I (the learning organism) can DO now, when what I am doing is not working to control my perceptions. I might alter the sign of some output link (go back instead of forward; stimulate the economy instead of depressing it), or I might do something that was not in the original repertoire at all (scream or bat my eyelids instead of going back or forward). This kind of reorganization is necessarily random. The organism has no way to determine what could lead to optimum behaviour in a smooth, gradient-search, kind of way.

>Learning is necessarily what Skinner calls 'superstitious'; it has nothing
>to do with truth, only with correlations.

I don't think anyone would dispute this, provided the word "correlations" is taken with generosity of spirit.

Enough. I hope this helps toward a reconciliation of your views and what PCT is about.

Martin

Date: Tue Feb 16, 1993 3:52 pm PST
Subject: My levels and Bob Clark's

[From Bill Powers (930216)] Bob Clark (9230205 and later) --

I wasn't accusing you of beginning with psychological constructs and then filling in lower-level systems. My point is different.

Some time between 1960, when we parted company, and 1973, when I published BCP, a change in my thinking about the levels seems to have occurred. Or maybe, being on my own, my direction of thought became clearer. This all seems to be clearer now that you're describing your hierarchical concepts once again.

At any rate, the "pre" idea was much like yours, that we were attempting to characterize human beings by identifying levels of control with various aspects of human functioning. Somewhere in that 13 years, I realized that this was not the right problem.

As I now think about it, the problem in understanding human nature is not so much to understand human beings as to understand the world that human beings experience. In this world I include not only the three-dimensional world around us, complete with living color, stereo sound, smellivision, and so forth, but also the "inner" world of imagination, memory, thought, reasoning, understanding

-- the whole world of inner commentary on sensory experience. In short, the world of experience includes everything experiencable, whether we think of it as being "inside" or "outside."

This world, to the best of my knowledge, originates in signals emitted into the nervous system by sensory receptors. That observation seems fundamental to me; to deny it would be to wreck the entire structure of physical theory, which I do not propose to do just yet. There is no way for the state of the world outside the nervous system to be registered in the brain without first appearing as a set of raw unanalyzed sensory signals. Nothing by way of information about the outside universe can get into the brain in any other way.

This means that the world we experience must consist of sensory signals and other signals derived from them. The "other signals derived from them" include the totality of what we can experience, from the taste of chocolate to Fermat's Last Theorem, as well as our experienced "interest" in that Theorem, if any, and any "thoughts" we may have about it. Nothing is exempt.

When I say "it's all perception" this is what I mean. We live inside a nervous system and all we know is what goes on inside that nervous system. Even our idea of the existence of the nervous system exists as a set of neural signals, perceptions. The physical world outside us is a network of hypotheses existing in neural networks in the brain. Part of this neural hypothesis is a conjecture to the effect that there is an objective physical world outside our sensors. Sciences like physics and chemistry are very well worked out neural hypotheses. At bottom, they rest on sensory experience and all that the brain can make of such experiences. Our very attribution of physical theory to objective phenomena is itself a phenomenon in the brain.

This changes the problem. Now the problem is to classify all of experience, not just experiences of other people. We may perceive another person driving a screw into a piece of wood as showing a "skill" type of control, but this leaves unexplained the screwdriver, the screw, the piece of wood, and the relations among them. Those are also perceptions, and they are being controlled. The term "skill" refers mainly to something about the person's organization, but to explain how a skill like that is carried out we have to explain the screw, screwdriver, wood, and relationship as well. The perceptual organization needed to represent these four things explains their existence for the actor; the actor's behavior is explained, in PCT, as control of these perceptions. Whether we characterize that control as constituting a "skill" is more or less beside the point. If we can explain the behavior in terms of controlling perceptions of wood, screw, and screwdriver individually, and in terms of adjusting those controlled perceptions to maintain control of a particular space-time relationship among them, we have explained "skill," too. But we have also explained how any person interacts with the world, whether the immediate world contains other living systems or not.

What I attempted to do with my definitions of levels was to represent the way the world seems to appear to us -- meaning, to myself as a representative human organism. This was very much an ideosyncratic first try, and it has undergone revisions as I have attempted to refine the descriptions. The process involved was quite unscientific, in that I didn't take any polls or do any objective experiments. I simply looked and listened and felt, and tried to understand what was going on from the standpoint that I was an observer watching the outputs of neural data-processing functions. "What am I taking for granted?", I asked over and over. What is it that I'm doing or experiencing that is so familiar and so self-evident that I don't even recognize it as a perception? What part of my experiences am I setting aside as having some special status, or treating as the background of more important things, or brushing out of the way so I can look at something more interesting?

The "relationship" level was a latecomer to the hierarchy. I had spent a lot of time looking for relationships between one perception and others, and between action and perception, but it took years for me to realize that relationship ITSELF is a perception. The same was true for all the levels added or modified since 1960. I had spoken for years about the "principles of control," without realizing that principles can't exist unless we perceive them, and to perceive them we necessarily have to have principle-perceiving functions. Similarly for "physics." What is physics, that I can know it exists? It's a perception, of course. If I couldn't perceive such a thing, it wouldn't exist for me. So what sort of thing is it? I have proposed calling such things "system concepts," for lack of any better term. And what other sorts of experiences are of that same sort? There are many, once you realize that this IS a sort of perception.

I think that the key to understanding how I think of the levels is to get into a mode of observation in which, as they say in Washington nowadays, "everything is on the table." No thought, no concept, no background perception, can be let go because it "doesn't count." Everything noticeable counts. Everything noticeable is evidence about what at least one brain is doing. If you accept the basic premise, that the experienced world begins as a set of unanalyzed sensory signals, the only conclusion is that everything noticeable is activity in a brain, and hence has to have a place made for it in a model of a brain.

I don't think that I've characterized the higher levels of perception very well. The most I hope to get across by the terms I use is the approach, the idea of calling into question everything we normally take for granted, all the operations and perceptions that we use in thinking about and acting on something ELSE. I don't think we'll arrive at a consensus on the levels until more people go through this very personal sort of exploration and report their findings.

Best, Bill P.

Date: Tue Feb 16, 1993 7:48 pm PST
Subject: semantics of control; technical definition of feedback

[Avery Andrews 930217.1314] Hans Blom, 930216

I don't think there is any mantric aspect to the BCP slogan, but there is a tricky semantic issue: engineers seem to use the word `control' to refer to pretty much any way in which a designer might get a machine to produce what a customer wants. But in a biological context, such as PCT, this won't do, since there are no designers or customers extrinsic to the systems themselves. What is needed is a notion of `control' that is intrinsic to the operations of living systems, & what PCT seems to mean by `control' is stabilization of something against disturbances, when this is surprising in terms of the gross physics of what is going on. E.g. it is not control when a tree stays upright against a moderate wind, since it is not surprising, given the physical nature of trees and the ground, and the way the former is typically stuck into the latter. But it is surprising when a person stays upright against a moderate wind, so that's control, leaving us with the problem of explaining how it happens.

`Perceptual control' (what feedback systems do) is then the basic means whereby control is achieved, tho there are many other auxiliary strategies. E.g. if you want to control for your bike to stay more or less in a given position, you can chain it to something, thereby exploiting the spring-like properties of wood, metal, concrete, etc. to resist disturbing forces (this is a kind of anticipation, or feedforward control, I guess).

In simple cases, you tell what's being controlled by applying disturbances, and seeing what aspects of the situation remain surprisingly unaffected. Your restaurant example is more complicated. Suppose somebody orders steak & is told there isn't any. If they're just controlling for having a nice dinner, they will just order something else, but if they're controlling for having steak, they'll go to another restaurant. But this doesn't fit the simple description I gave above, since what is being controlled is not a simple continuously varying quantity that has a value at each point in time, but a sort of proposition `we have a nice dinner/a steak for dinner tonight'. Perhaps one could say that in spite of the successive revelation of information that would seem to make the achievement of the goal less likely, the goal gets achieved anyway. So what the people would actually be controlling for is the apparent likelihood of getting a dinner/steak at a reasonable time in the evening (controlling for apparent likelihood involves a lot of internal simulation, maybe reasoning, etc,).

Re Kuo: I don't see how to get the `technical definition' of feedback out of what Kuo says. Here's an actual quote of the doctrine:

> `technically, a *feedback* system is one in which an error signal
> directly drives a corrective adjustment at the site where the error
> is introduced' (Abbs & Winstein 1990 `Functional Contributions of
> Rapid and Automatic Sensory-based Adjustments to Motor Output',
> Attention and Performance XIII, pp. 627-652.

This appears in a section entitled `Contributions outside the closed-loop model', whose wording makes it quite clear that the authors think that the engineering conception of feedback is of very limited use, and that the lip-aperture adjustments they studied go beyond it (since one lip is prevented from closing, and the other moves further to make for it).

There is perhaps a way they could get their doctrine by looking at control theory pictures but not thinking hard enough about them. In the intro diagrams there is typically something you can point at (the Plant), with a part (such as a driveshaft) with a property (rpms) that can be measured, and is what the customer wants, and is in some sense what the Plant was built to produce. So the control system monitors the desired property of the Plant (the `output', on the usual labelling), and injects control signals to stabilize it at a desired level. (Aside from the throttle, the other common intro picture seems to be the fermentation vat, where the Plant aspects that are being controlled are concentrations of various substances.)

But in the biological examples, although these components are in a sense there from a functional point of view, they don't correspond to discrete physical things that can be pointed to. So if a lizard is regulating its body temperature by scuttling between sun and shade, it's pretty hard to say what the

`Plant' is, except perhaps that it's the whole world, including the sun. And the `output' of this plant is just the body temperature of the lizard, as perceived by whatever sensor the lizard employs.

In the more general case, animals have muscles that affect how things are, and sensors that detect how things are, but the connections between what is done in a physical sense and the biologically relevant properties of the situation that are affected are quite complex, so that one can't regard the creatures effectors as a plant with a determinate aspect to be its `output'. I guess you could in fact regard the Powers diagram as treating the entire environment as the Plant, including, in the case of higher level control systems, the lower-level ones, so that, in a sense, he's turned the conventional diagrams inside out, with the Plant surrounding the control system, rather than being a little box with the bits and pieces of the control system surrounding it.

To me it seems that even the usual thermostat example goes beyond the bounds of the supposed `technical definition' - if one regards the plant as being the furnace in the basement. For `errors' (drops in temperature in the vicinity of the thermostat) can be introduced all over the place (by opening doors and windows), so what is controlled isn't a `direct' output of the Plant (my wording), and neither are errors `corrected at the site where they are introduced' (Abbs & Winstein's wording) - they're introduced at, say, the front door, but corrected in the basement.

Engineers seem to draw their pictures in a fairly freewheeling manner, as aids to thought rather than attempted limitations to its scope, so my guess is that the form of the pictures has had a sinister effect on the minds of psychologists, and has been interpreted in ways that don't actually make any sense, and impose invalid limitations on the scope of closed-loop systems. Note that A&W don't just proclaim a technical definition of feedback, but claim that it indicates a fundamental limitation of closed loop systems: they begin the section (Contributions outside the closed-loop model) with the words:

> `Often, available formalisms developed for designing man-made systems offer substantial insights into biological organisms. In other cases, these constructs, carried too far, limit the paradigms and hypotheses that would otherwise be naturally pursued.'

My guess is that the limitation resides more in the minds of the psychologists (including Schmidt, who says something similar to A&W in his motor behavior textbook) than in the actual engineering concept, in this case.

The Kuo book looks nice - in both the editions I looked at, by the way, (1967, 1987), it says that `humans are probably the most complicated and sophisticated feedback control system in existence'. Sadly, my complex analysis and differential equations are too feeble for me to be able to just read it.

Avery.Andrews@anu.edu.au

Date: Tue Feb 16, 1993 8:26 pm PST
Subject: learning, place cells

Gary,

Why doesn't the news network post posts in strict order--Rick's reply to Martin's reply came before Martin's reply (which in turn were replies to Rick's reply to my post on learning).

Martin,
You stated that you identified about a dozen different learning paths--could you send me a list of these paths--I would be very interested.

Rick,
Thanks for clarifying that PCT doesn't say much about the learning part of operant conditioning. I imagine one of Martin's 12 learning ways will account for operant forms of learning and another for classical forms of learning.

Anyone interested,
A few days ago I sent a post on place cells and place fields. To continue with that and give a brief summary of where I was going with it: I began with the thought that the one region of space which in normal everyday reality is not represented within the place fields of any place cells is Within the Space of One's Own Body (at least above the shoulders--a minor point to argue about). This is because there is no functional need to represent this region of space since we obviously don't have to be concerned about maneuvering around it or misplacing it (As Buckaroo Bonzai says, "Wherever you go, there you are.")

But place fields do drift and they do so for functional reasons. But whether for functional or not, what would happen if the place fields of these cells drifted over into that region of space occupied by our bodies? I suggest that one would have an "out-of-body" experience.

Now before anyone thinks this is just too weird, let me suggest that what I am doing with this idea is putting "out of body" reality on the same epistemological plane as our "everyday" reality. (Nothings is being said here about Boss Reality). Since, as we know, reality is constructed by us and, I would argue, is done so partly by the firing of place cells being Equivelant to the experience of object x THERE or THERE, then there is no leap to be made in stating that an out of body experience occurs when cells concerned with body space fire, giving the organism the same ability to interact/manipulate the environment which includes self space, that it would have to interact/manipulate the environment that (doesn't normally) include selfspace.

As far as I can tell, this is an original perspective on such an issue in the sense that it doesn't (1) simply relegate out of body stuff to "spiritual matters" or (2) simply attempt to EXPLAIN AWAY the phenomenon as somehow abnormal or subreal on the basis of faulty chemical composition or faulty processing. Instead it accounts for the apparent realness of the experience, putting it on the same plane with the apparent realness of our normal experience.

(Such an idea reminds me of that one famous painting whose name and artist I don't remember, in which Plato is pointing up and Aristotle is pointing down, each pointing to "Where it's AT" in their opinion. That the "spiritual" is down here with everything else is the implication of this idea that both Aristotle and I would/do really get off on.

My purpose here is not to talk spirituality, of course, but to bring up the Neuropsychology of it (part of it) and more importantly to say something about epistemology. The next question is "Why might these place fields move to cover selfspace--what is the functional advantage of being able to watch oneself perform?" In case I am not being clear, these experiences I am refering to happen in various occasions--some say flying while dreaming is an example, certainly this happens to athletes when they are going to make "an impossible catch" or to drivers when they are about to go through the windshield of their car....

Though I made no remarks about PCT in this, I don't think it is irrelevant to PCT. But I probably won't bring it up again in the near future unless someone comments.

Mark

Date: Wed Feb 17, 1993 5:53 am PST
Subject: CL Research Reports; Graphics Library

From Greg Williams (930217)

1. Most CSG members will have received the Winter '93 issue of CLOSED LOOP by now. It inaugurates inclusion of "Research Reports" with two papers. I'd like to encourage submission of articles for possible inclusion in future issues (send manuscripts to me at 460 Black Lick Rd., Gravel Switch, KY 40328 U.S.A. or via e-mail), AND I'd like to have some volunteers to read/evaluate/critique such submissions. Perhaps we could set up an ongoing panel of reviewers? Publishing semi-peer-reviewed Research Reports in CLOSED LOOP is another step toward a JOURNAL OF LIVING CONTROL SYSTEMS.

2. Fast, cheap, and NOT out of control graphics library for Turbo C and MS C: Michael Jones, GRAPHICS PROGRAMMING POWERPACK, Sams Publishing, 11711 North College, Carmel, IN 46032, 1992, $24.95, ISBN 0-672-30120-2. In many bookstores now. Comes with a 3.5" disk with lots of examples and small, medium, large libraries. No run-time fees, EGA/VGA support, VERY fast, XOR works with everything, even text (which is available in three nicely sized fonts). Ellipses are ca. 30% faster than BGI; rectangles ca. 50% faster. Mouse support routines, also. Lots of PCX manipulation stuff. Some sound stuff. Our graphics library of choice for now. A subset of the Genus Microprogramming GX libraries at a bargain-basement price.

A minority member of the net (I control for line width ca. 75 chars.),

Greg

Date: Wed Feb 17, 1993 6:08 am PST
Subject: Focus Control; Taylor

[from Gary Cziko 930217.0343 GMT]

Martin (930216 16:20) in his illuminating reply to Hans Blom (which reminds me how far many of us on CSGnet have come in our understanding of PCT) says:

>If you can predict exactly what the results of actions will be, you don't
>need control. If you can't predict at all how your actions will affect
>the world, you can't perform control. If pushing and pulling are both

>equally likely to bring the object closer or further, ...

When I show photographic slides ("diapositives" for our European friends) to my family and friends, some are initially out of focus. I have never been able to determine from the perceived blurriness which way to turn the focusing knob to bring it into focus. I try one direction. If it gets better I keep going. If it gets worse, I turn the other way. And I can never tell when I've gotten the image in optimal focus until I go past the best setting and then I have to back track.

> and it is going to move the same way even if you do nothing,..

And sometimes while viewing the slide it "pops" (bulges) out of focus because of the heating it undergoes in front of the projector bulb. It usually pops in one direction, but if the slide is in backwards (sometimes by error, sometimes by intent), it will pop in the other direction.

>you can't bring the distance to where you want it.

And yet I have no real problem controlling the focus of my slides. So,
Martin, either control systems are more robust than you think or I am not understanding what you mean here and my slide example is not relevant. Which is it?--Gary
Date: Wed Feb 17, 1993 8:55 am PST
Subject: my levels, semantics of control

[From Rick Marken (930217.0800)] Bill Powers (930216) --

Wow!! Great post. Nice to have you back on the net!

I still think there might be a way at explore this phenomenon (the nature of human experience) "objectively" using something like the perceptual study described in my "Hierarchical Behavior of Perception" paper.

Avery Andrews (930217.1314) to Hans Blom (930216) --

Excellent description of problems with the term "control". Maybe you COULD write the PCT based paper that will be accepted by BBS or Psycholoquy. Very nice post.

Best regards Rick

Date: Wed Feb 17, 1993 9:44 am PST
Subject: Re: Focus Control; Taylor

[Martin Taylor 930217 12:10] Gary Cziko 930217.0343

>>If you can predict exactly what the results of actions will be, you don't
>>need control. If you can't predict at all how your actions will affect
>>the world, you can't perform control. If pushing and pulling are both
>>equally likely to bring the object closer or further, ...
>And yet I have no real problem controlling the focus of my slides. So,
>Martin, either control systems are more robust than you think or I am not
>understanding what you mean here and my slide example is not relevant.
>Whic is it?--Gary

The slide example is not relevant, except as an instance in which partial prediction is possible. If the focus varied in an unknown direction by an unknown amount every millisecond, it would count as a system in which you can't predict the effect of your action. As it is, you move the focus knob, and except for perhaps one or two isolated moments, you predict (probably correctly) that the focus is changing in a more or less uniform (at least monotonic) manner. The only thing you don't initially know is which way it is changing, but you can see that pretty quickly. If you leave the focus alone, it doesn't change, except for one or two isolated moments. Pretty good predictability.

Martin

Date: Wed Feb 17, 1993 10:47 am PST
Subject: Re: re: re: re: control article

[Hans Blom, 930217] Rick Marken (930216.1200)

> "Control of perception" is neither
>a mantra nore a philosophy -- it is a fact about control system operation;

I have stricken the word 'fact' from my vocabulary. As Bill Powers would say, "it's all perception". If this perception works for you, fine. In many cases, it works for me, too. If I sometimes need a different point of view, I hope I'm entitled to it, as well.

>If you think that this observation is incorrect; if you have a way of
>showing that it is not always the functional equivalent of the perceptual
>signal in a control loop that is controlled -- that sometimes, for
>example, it is the functional equivalent of the output variable that is
>controlled -- then we would REALLY like to know about it; I would like to
>know about it because I have spent a good deal of my life trying to show
>that behavior is the control of perception; if it's sometimes the control
>of output or error or disturbance or some other variable in the loop the
>I'm going to have some 'splainin' to do.

I do not think your observation is incorrect, I just hate the word 'always'. I agree with BP's "it's all perception" in the sense that perceptions (of the outside world and of our inner physical and mental mechanisms) are the only sources of information available to us. But perceptions are built upon and result in higher level things that I would not call perceptions anymore. Beliefs, superstitions, the 'facts' of our lives. All those together constitute what I call a model (of the world, ourselves included). A model is, technically, always a simplification, and always has a purpose. That it is a simplification is due to the facts that we have experienced only a limited set of perceptions, and that our processing of those perceptions must be done by a mere three pounds of flesh. Models are never unique; it is always possible to translate one model into another, equivalent one. Sometimes a simple, approximate model works well enough, sometimes only a very complex and very accurate one will do, depending upon the goal that it serves. The highest purpose of the biological model is, in my opinion, best described by Dawkins: transmission of genes. Everything serves that supreme goal. The evolutionary process has weeded out every organism that did not serve its purpose well enough. A high degree of optimization has taken place during billions of years, and in that sense all currently existing organisms can surely be called well- designed control systems. Control systems, because they need to achieve a goal. There are numerous ways to achieve that goal. Viruses, bacteria, cats and humans do it differently, thus far equally successfully. All other goals are sub-goals, designed through evolution to serve the one supreme goal. The sub-goals of each organism are uniquely related with its poten- tial for actions, i.e. its body. A virus needs very few perceptions to achieve its goal; it mainly relies on the forces of nature ('free' energy) to work for it. A human, on the other hand, cannot survive without a great many perceptions.

In short, I think that your perspective is extremely valuable when you study human behavior. A different perspective might be better for me, because I study very simple things like control systems. I sympathize with your "I have spent a good deal of my life trying to show that behavior is the control of perception". I, also, have spent (a good deal of) my life acquiring my own personal perspective. Let's by all means keep exchanging perspectives! Sometimes it seems less limiting to have two different perspectives on the same reality at the same time. Could that be why binocular vision proved to be successful?

(Martin Taylor 930216 16:20)

>> where the system's gain is low. Sometimes that won't matter, but in
>>high-performance systems it DOES matter. Then it is necessary to know (a
>>priori or through learning/adaptation) and account for the nature of the
>>non-linarity in the control process.

>The only problem I see in this is the word "know" in the second-last line.
>...   Usually, "is able to do" is a reasonable
>substitute for "know how to do," in cases where the distinction matters.

In a control system, "to know" means "to have available", either in the form of models or model parameters that compactly store (relations between) previous perceptions [example: what is this patient's sensitivity for drug X?], or in the form of potentials for action, different means to achieve some goal through action [example: I could apply either drug Y or drug Z]. In the text above, I referred to the former.

>>The fundamental point of control theory is "if you can't predict it, you
>>can't do anything about it".

>If you can predict exactly what the results of actions will be, you don't
>need control. If you can't predict at all how your actions will affect
>the world, you can't perform control.

Control engineers have a broader conception of control than you seem to do. Control does not necessarily imply feedback. In fact, engineers prefer nonfeedback systems if at all possible, because they cannot possibly have stability problems. Regrettably, non-feedback control is possible only if the system to be controlled in invariable and not significantly subject to disturbances. An example of a non-feedback control system: the direction of a car's steering wheel controls the direction of the front wheels [Kuo, page 2]. Let me therefore modify your first statement as follows:

> If you can predict exactly what the results of actions will be, you
> don't need feedback.

Your second statement is essentially correct. That may not be the end, however. An adaptive control system can learn the relations between its actions and their effects by generating an action and observing its effect. If a non-zero effect results, control becomes possible.

>Avoidance is that the ECS brings its percept to a value different from
>the reference in either direction.

Can I equate "a value different from the reference" with "a new reference"? If so, I would not call it avoidance.

>In answering Rick (930211.0800), you get into questions of decision.
>There are at least three aspects to decision. And again it relates to
>"overt" and "covert" operation.

>(1) If two actions are incompatible in the world, the control loops that
>involve them are mutually inhibitory. Only one of the relevant perceptions
>(or neither) will actually be controlled. We have had some discussion of
>this in the past. I like to look on it as a question of loop impedance,
>in which the coupled real-world systems affect each other's impedance and
>hence the gain of the loops in which they participate.

The term "impedance" assumes a linear and static world. What if the world kicks back, as in hysteresis?

>(2) The second aspect of decision is in the distribution of output from an
>ECS. It is conceptually possible for an output function to alter the
>distribution of its outputs as a function of the error. That could be seen
>as decision (I am far in space and near in time to somewhere I want to
>perceive myself, so I take the car, whereas if I were near in space and
>far in time, I would walk).

I would call this a selection between different possibilities of action.

>(3) A third aspect of decision is in the program level. As I see it, a
>program is a set of bifurcations (or splits of higher multiplicity). Each
>bifurcation selects a reference sequence for lower ECSs. That, too, is
>decision.

Again, I prefer the word "selection".

> But from the viewpoint of the organism
>that uses the loop, things are quite different. The perceptual signal
>is your only way to determine what is going on in the world. You can't
>detect your outputs, and you can't control them. What you can control
>is the perceptions that result from your outputs, and that inherently
>affects the outputs themselves.

Luckily, humans are wired in such a way that they can sense their outputs; this is called the "body image". Control systems, too, know (use) their current (and sometimes previous) outputs. These, after all, will generally have an effect on things yet to come.

>> Once something is learnt, it can be used to predict ('what

>>happens if I do X') or to control ('what should I do to achieve Y'). In
>>this view, learning or 'reorganization' is not a random process, but a
>>process much like curve-fitting in statistics.

>There are different kinds of learning within the PCT framework. Some are
>necessarily random, some not.

In adaptive control theory there are two kinds of learning as well. One is passive learning, through observation of the correlation between input and output during normal control. This works only if there is something to correlate, not if both input and output remain constant all the time. In such cases, even though control may be perfect, nothing is learned. For learning to occur, the input must be "persistently exciting", e.g. due to a sufficient noise level in the feedback loop.

The other type of learning is active, through addition of some kind of test signal to the input signal that would be required for normal control. This test signal may be white noise, a pseudo-random binary sequence, or some more intelligently chosen signal. A more intelligent choice, by the way, is possible only if some details of the model have already been established. Active exploration helps in establishing a better model sooner, but if overdone deteriorates the quality of the control too much. It is often hard to find a good compromise between control and exploration, but it is necessary to do so in a changing world, i.e. when the system to be controlled may change over time.

>Enough. I hope this helps toward a reconciliation of your views and what
>PCT is about.

I think that by now I understand what PCT is about. I have followed and enjoyed the discussions for more than a year now, mostly quietly. Once in a while I grab the chance to vent some of my ideas, which are more or less related, hoping for a useful reply -- usually not in vain. Reconciliation is not what I look for; I find that friction -- clashing points of view -- generates much more creative energy.

(Avery Andrews 930217.1314)

>I don't think there is any mantric aspect to the BCP slogan, but
>there is a tricky semantic issue: engineers seem to use the word
>`control' to refer to pretty much any way in which a designer might
>get a machine to produce what a customer wants. But in a biological
>context, such as PCT, this won't do, since there are no designers or
>customers extrinsic to the systems themselves.

Don't think too highly of designers! Many of their best ideas seem to be the result of chance -- called good luck -- as well. It is just that the bad ideas are soon discarded, usually...

>Re Kuo: I don't see how to get the `technical definition' of feedback
>out of what Kuo says. Here's an actual quote of the doctrine:

> `technically, a *feedback* system is one in which an error signal
> directly drives a corrective adjustment at the site where the error
> is introduced' (Abbs & Winstein 1990 `Functional Contributions of
> Rapid and Automatic Sensory-based Adjustments to Motor Output',

> Attention and Performance XIII, pp. 627-652.

>This appears in a section entitled `Contributions outside the
>closed-loop model', whose wording makes it quite clear that the authors
>think that the engineering conception of feedback is of very limited
>use, and that the lip-aperture adjustments they studied go beyond it
>(since one lip is prevented from closing, and the other moves further
>to make for it). ... Note that A&W don't just proclaim a tech-
>nical definition of feedback, but claim that it indicates a fundamental
>limitation of closed loop systems: they begin the section (Contributions
>outside the closed-loop model) with the words:

> `Often, available formalisms developed for designing man-made systems
> offer substantial insights into biological organisms. In other cases,
> these constructs, carried too far, limit the paradigms and hypotheses
> that would otherwise be naturally pursued.'

>My guess is that the limitation resides more in the minds of the
>psychologists (including Schmidt, who says something similar to
>A&W in his motor behavior textbook) than in the actual engineering
>concept, in this case.

Engineers and psychologists are not close neighbors. They speak different languages, have a different culture and work on different problems, although it is fascinating to discover similarities. I believe that engineers can learn as much from psychologists as the other way around. Doesn't this list show it?

I would modify A&W's definition as follows:

a feedback system is one in which an error signal provides information about how to perform a corrective adjustment

The error signal need not 'directly drive' a correction (its square or integral or some such might be used). The 'at the site where the error is introduced' is not essential, either.

>The Kuo book looks nice - in both the editions I looked at, by the way,
>(1967, 1987), it says that `humans are probably the most complicated
>and sophisticated feedback control system in existence'.

How about Gaia?

Date: Wed Feb 17, 1993 11:54 am PST
Subject: meaning of control; how smart is Gaia

[Avery.Andrews 930218.0600] (Hans Blom, 930217)

>>If you can predict exactly what the results of actions will be, you don't
>>need control. If you can't predict at all how your actions will affect
>>the world, you can't perform control.
>
>Control engineers have a broader conception of control than you seem to do.

Semantics strikes again. One palliative might be to refer to our kind of control as `perceptual control' (at least in foundational discussions), and stop trying to load a special technical meaning onto the unmodified word `control'. Aside from the fact that it doesn't have this meaning in the sider community we would like to interact with, it is also a fact that organisms do use techniques other than P.C. to stabilize aspects of the environment that are important to them.

>I would modify A&W's definition as follows:
>
> a feedback system is one in which an error signal provides infor-
> mation about how to perform a corrective adjustment
>
>The error signal need not 'directly drive' a correction (its square or
>integral or some such might be used). The 'at the site where the error is
>introduced' is not essential, either.

This is good.

>Engineers and psychologists are not close neighbors. They speak different
>languages, have a different culture and work on different problems,
>although it is fascinating to discover similarities. I believe that
>engineers can learn as much from psychologists as the other way around.
>Doesn't this list show it?

Yes it does. I was being nasty about the psychologists because, for some reason, people who make up limitations that don't actually exist annoy me.

>[re human being as most sophisticated control system] How about Gaia?

Probably not, unless she's clever enough to get rid of us without zapping everything else.

Avery.Andrews@anu.edu.au

Date: Wed Feb 17, 1993 12:08 pm PST
Subject: Re: control article, focus

[From Rick Marken (930217.1100)]

I said:
> "Control of perception" is neither
>a mantra nore a philosophy -- it is a fact about control system operation;

Hans Blom (930217) replies --

>I have stricken the word 'fact' from my vocabulary. As Bill Powers would
>say, "it's all perception". If this perception works for you, fine. In many
>cases, it works for me, too. If I sometimes need a different point of view,
>I hope I'm entitled to it, as well.

I give. I thought you might have some EVIDENCE for your point of view. If not, then just enjoy it! You are certainly "entitled" to it.

Martin Taylor (930217 12:10) to Gary Cziko (930217.0343)

>The slide example is not relevant, except as an instance in which partial
> prediction is possible.

I think the slide is an excellent example of control with zero predictability of the effect of the control action; another easy to demonstrate example is tuning a radio. There is no way to predict (better then 50%) whether a clockwise or counterclockwise turn of the tuning dial will improve or degrade the degree of tune -- you just have to turn and see what happens (well, with the radio you could look at the dial and see which way you "should" turn to get it tuned -- but let's eliminate this; we can tune a radio while driving, say, so all we have on which to base our control actions is the behavior of the controlled variable itself (the degree of "tunedness" or "focusedness")). Focusing and tuning just show that there are certain kinds of control that CANNOT be done by a single level control system. A "tuning" control system that, say, turned the radio dial in a certain direction as a function of error would only be able to tune radios that happens to be off tune in the "right" direction. We can tune radios because we can perceive and (hence) control whether we are "approaching" or "moving away" from tuned (focused). This is a higher level variable then "degree of tune" (or focus). It is controlled (apparently) by changing the polarity of the relationship between error and output in the "degree of tune" control system.

Best Rick

Date: Wed Feb 17, 1993 2:38 pm PST
Subject: Catching up on many subjects

[From Bill Powers (930217.1030)]

Had a nice visit with Rick Marken, but as he mentioned it was frustrating. That Mac is a very hard computer to program. The amount of foreplay required just to read a keystroke is unbelievable. Nice to see Dag Forssell and Christine, too.

Martin Taylor (930207.1320) --

RE: Flinging suitcases -

I think this phenomenon is taking on the dimensions of a myth. In fact it hardly ever happens. Usually you just pick up the suitcase, whether its contents weigh 3 pounds or 15 pounds (anywhere in the usual range of weights). You might be surprised at how heavy or light it is, but this doesn't cause you to drop it or fling it. You judge its weight by sensing the effort made when the position control systems raise the suitcase off the floor. The effort adjusts quite automatically to the weight.

As to reaching in the wrong place for the baking soda even though you can see it in a new place, I think this shows that the initial reaching is probably under kinesthetic control. You reach, then look for the object to do the final grasping. But we're all different; I don't think we can draw any general conclusions about how human beings reach for the baking soda.

Most of these anecdotes call for systematic experimental test. If people do use programs for some actions, we have to test to see what aspects of them are disturbance-resistant. Just watching something happen doesn't tell us a great deal about what's controlled.

Avery.Andrews (930208.0914) --

>... the limitations of feedback control of errors that have
>already occurred are a familiar theme in the anti-feedback literature.

That theme is based mainly on thinking of error qualitatively: either there's an error, or there isn't. This kind of either-or thinking can't even explain a simple artificial control system.

Oded Maler (930206.1000) --

>I don't know if that is what you meant, but I think it's about
>time that you teach your Little Man to draw and write. A motor
>program that draws "A"s or triangles will be a very impressive
>demonstration of "your" notion of a motor-program, especially
>if the size, the location (relative to the shoulder) can also
>be adjusted by higher references, and there is a lot of noise
>between the levels.

The Little Man can already do this. By sending a sine and cosine wave to the reference inputs of the x and y visual control systems, the program can make the fingertip draw a circle continuously around the target, whether the target is stationary or moving. By specifying more complex waveforms one could produce any pattern of movements relative to the target. I don't know how to model a perceptual system that can perceive the patterns of writing, so I can't build the level of control that produces those waveforms in the manner of a control system. But the model provides for reference signal inputs that will make the fingertip movements follow any prescribed pattern.

There isn't any "noise between levels" in this system. Why should such noise exist?

RE: aboutness and intentionality

>So what is the perceptual signal about?

Nobody really knows. My conjecture is that there is a regular lawful universe outside the nervous system that constrains the relationships between the perceptions I call "my actions" and the perceptions I call "the effects of my actions." For an extensive and orderly system of hypotheses about the nature of this external universe, consult any physics or chemistry textbook.

Unfortunately, perceptual signals are not direct representations of that external universe, but many-to-one functions of neural signals that themselves ambiguously reflect interactions of sensory nerve-endings with local conditions outside the nervous system (for example, interactions between visual receptors and incident radiation -- hypothetical photons).

>But this metaphor is the convention underlying the the use of
>language as a medium for communicating meaning.

The PCT view, I think, is that meaning is not communicated but evoked. What I say evokes experiences in you drawn from your own past interactions with your world. Those experiences are the meanings that my words have for you. To the extent that we have had similar experiences, my words may evoke experiences in you that are somewhat like the experiences I attach to those same words (and linguistic forms). Determining the extent of the match between my meanings and yours is an unsolved problem, perhaps unsolvable. Our apparent ability to communicate to our mutual satisfaction is unexplained, and perhaps not as well-developed as it may seem. If we could truly transmit meanings from one mind to another in any simple way, most of what goes on on this net would be unnecessary.

(you to Rick)

>The signal in my home theromstat is about the temperature in my
>house, and its value in this signal, represents some fraction
>of the set of all possible worlds.

As a disembodied third-party observer, you can say this. As a thermostat, you could not say it. To the thermostat, the position of the movable contact IS the temperature. In the thermostat's perceptual world, the agitated molecules in the air do not exist. You can speak of what a perceptual signal is "about" only if you can perceive both the signal and the external thing to which the signal corresponds. That is something that only a disembodied third-party observer can do, and then only for some other system. The third-party observer (TPO) can't compare his or her perceptions of the other system and its environment with the actual condition of the other system and its environment. The only way the TPO can know of either the other system or its environment is in the form of perceptual signals inside the TPO. Unless you want to introduce an infinite regression of TPOs, or grant the TPO some nonhuman or extrasensory way of knowing about the environment, you therefore can't give an authoritative account of what any perception is "about."

Hans Blom (930208) --

>> Human beings hardly ever control the "full trajectory."

>If that is the case, 'new types of control', which do not
>try to maintain minimum error between a reference value and
>perceptions at all times, might provide superior performance
>in some cases. Or greater ease. When I fly to New York, I
>(attempt to) control my destination, but in the plane I have
>to trust the pilot. Part of my trajectory will be, as far as I
>am concerned, ballistic.

I think you're going about this backward. When we study human behavior, we aren't comparing it with some "optimal" or "best" way of controlling. We're just trying to understand what people are actually controlling under various circumstances. In some regards, people control things very well indeed, by clever means that surpass what any engineer knows how to build. In other ways, people control stupidly and poorly, and suffer the consequences.

Most to the point, people use the means available to achieve whatever degree of control is possible. When I buy a ticket on an airplane, show up for the flight, and strap myself in, I have done all that is possible to get myself to the destination by that means of transport. So that's all the control I have; if the plane is highjacked to another destination, that disturbance is beyond my ability to resist. All I can do is wait until the plane lands and I can get off it, and then start controlling again for getting to the destination by some other means. It could easily be that I would have arrived at the destination sooner, even without the highjacking, by taking a bus. But I didn't think of that. People are not optimal controllers; they just do the best they can.

>_The_ controlled variable? What makes control in organisms
>so difficult to study is the simultaneity of a great many
>different ongoing goals, whose importance may, moreover,
>fluctuate from moment to moment due to influences beyond our
>control and usually beyond our knowledge.

The hierarchical model helps here, because higher-level goals change more slowly than lower-level goals. Many of the fluctuations in conditions are just disturbances, which lower- level systems automatically compensate for by adjusting lower- level goals. Much of the apparently chaotic nature of behavior becomes more understandable when we ask about higher-leel goals. We can then understand many external events as disturbances, and see how the changes in detailed behavior oppose their effects. This reveals regularity where formerly we couldn't see any. I think that most behavior is actually quite regular, once we understand what's being controlled at many levels.

You're right about the fact that more variables are under control than we can measure in any one experiment. But it's interesting that without much trouble we can get those other variables to remain constant enough to get good repeatable data.

>To me, this seems to be a clearcut case: a high jumper
>wants to jump as high as possible, period. An objective
>measure is provided to test that performance. All else is unimportant ...

The highest-level goal is to win the contest, not to jump as high as possible. There is strategy involved as well as just trying to produce maximum effort. Some jumpers will pass at a certain height, saving their strength for later: they don't try to jump at all. Also, if you assume that every time you see a high jumper the objective is to jump as high as possible, you will usually be wrong; most of the time, the high jumper is just trying to go high enough to clear the bar. On many other occasions, the jumper may not be concerned at all with controlling for height. The jumper might be working on the approach or the takeoff, or the form at the peak of the trajectory, or the flip that raises the legs at the critical instant, and not be worrying at all about maximum height.

You can't tell what a person is doing just by looking at what the person is doing. The test for the controlled variable helps you to understand what is actually being controlled (as opposed to what you logically assume is being controlled).

>Is a case where 'tendons and muscles can be ripped loose'
>really an indication of 'an intact set of control systems'?
>I consider that to be pathology, a control system gone
>haywire, operating beyond its design limits.

Certainly it is. If pathology is involved, it is a higher-level system that is misusing its lower-level control systems. Is it pathological for a father to lift a Volkswagen off his child, suffering torn muscles and ligaments (and a lot of pain) as a result? When a person shoots himself in the head, all the control systems for grasping the gun, aiming it, and pulling the trigger are working perfectly well until the last moment; all that's haywire is the higher-level system that has chosen this outcome. And even that choice may not be pathological, if the person is facing torture or the pain and humiliation of a vicious disease by staying alive.

>I would maintain that one of the most important of an
>organism's objectives is, at all times, not to seriously damage itself.

Normally, perhaps. Not always.

>> Even the muscles work differently from the servo motors
>> that engineers use. They don't apply forces directly, but
>> by shortening the contractile elements in the muscle to
>> alter the resting length of the series spring component.

>That is also what Levine and Loeb maintain, and they show
>how difficult it is to reach top performance with such
>'difficult' actuators.

They are not difficult actuators; they are marvelously effective and efficient. In order to sustain a constant load, a motor must generate a continuous stalled torque at tremendous energy cost. A muscle simply twitches some of its contractile fibers to the short position, stretching the spring component which then sustains the required force. In principle no energy at all need be expended beyond the initial amount needed to stretch the spring. In practice, the contractile element relaxes over some time after the initial shortening (returning about 40% of the expended energy to chemical form) so it must be twitched again and again, lowering the efficiency somewhat.

Such actuators are "difficult" for a designer only when the wrong design is used. The spinal reflexes make beautiful use of these actuators. Without feedback, of course, they would be difficult indeed.

>> The reason a human being can't perform a mathematically
>> optimal jump is simply the rocket problem: you would need
>> to produce an impulse of muscle force of zero duration and
>> infinite amplitude. That would hardly be a feasible
>> solution for a servomechanism, either.

>Impulses are not required, step functions will do nicely.
>After all, a trainer just want to study the peak performance
>that a real individual is capable of given her motor
>equipment, and search for whatever means there are to teach her
>to fire her nerves in such a way that this peak performance is
>reached.

A step function would be the most inefficient method of propulsion, because after it is turned on, it must continually counteract the force of gravity. You could run out of fuel without ever actually lifting anything off the ground. The most efficient use of a finite energy store in achieving a high trajectory is to use it all at once in the first instant, converting it to upward momentum. This is why weight-lifters do a "clean-and-jerk", lifting the weight to a new support position in a single burst of maximum force.

There is no way to "fire her nerves in such a way that peak performance is reached." How the nerves fire must be adjusted for disturbances and changes in system parameters; the outcome, not the output, is under control.

>Levine and Loeb do not say that feedback is too slow; bang-
>bang control requires very accurate timing. They say that when
>the need for performance becomes extreme, protection
>mechanisms are requir- ed to prevent muscles and tendons from
>being torn loose.

I suspect that this is another of those myths about control, this time about spinal control systems. For a long time, it was thought that the tendon reflex had the purpose of "limiting" muscle tension to prevent damage. This was based on the erroneous finding that the Golgi (tendon) receptors responded only to large forces. In fact, these receptors start firing with tensions in the tendon as small as 0.1 gram (the maximum possible force due to the biceps is something like 800 kilograms). I quote from

McMahon, Thomas A., _Muscles, reflexes, and locomotion_
(Princeton: ;Princeton University Press, 1984).

"It was supposed originally that the tendon organs did nothing until safe muscle loads were exceeded, but later evidence showed that tendon organs respond to less than 0.1 g of force applied directly to the base of the capsule (Houk et. al., 1971)" (p. 149).

In fact, nothing prevents damage to muscles but the fact that higher systems do not normally send large enough reference signals to the spinal control loops to strain the muscles. When necessary, they are perfectly capable of doing so.

As to the "accurate timing" of bang-bang control systems, how is this timing adjusted in the presence of disturbances to maintain the result in the same form? I don't believe there are really any bang-bang control systems in the human body. From what you say about Levine and Loeb's proposals, I don't think they know much about control systems, living or otherwise. They seem to be relying on outdated information about the tendon reflex, at least. What say those PCTers who have read their stuff?

>Feedback from those protective sensors would probably be too
>slow if train- ing did not slowly familiarize the high jumper
>with the sensations that they provide (1). This is much like
>walking as closely to the abyss as you dare without risking
>the damage that a fall would cause (2). The fall would
>provide you with feedback, of course, but you wouldn't want
>_that_ feedback, would you?

Since the receptors aren't "protective" in the first place, but simply provide feedback proportional to the load, none of that means anything.

>> Human beings hardly ever control the "full trajectory."
>> They control the variables that matter to them.

>Yes. And bodily (and mental) integrity matters a great deal.

I disagree. This is like saying that organisms control for "survival." Organisms control specific variables relative to specific adjustable reference levels. An outcome of doing so may be that the organism "survives" or preserves "physical and mental integrity," but that is not a concern of the organism. It's an opinion of a third-party observer. I don't think there is any reference signal specifying survival or integrity. Organisms don't survive or preserve their integrity anyway. They all die.

>> "Stabilizing control" is something of a misnomer,
>> suggesting that all that a control system does is to keep
>> something constant. More generally, it makes the perceptual
>> signal track the reference signal.

>Exactly how would you know that the jumper follows a
>reference signal when for the very first time she jumps higher
>than she ever did before? How does the reference signal get
>established in the first place?

The trajectory is a side-effect of controlling variables that the jumper can control. It is not itself a controlled variable. Once the jumper has left the ground, there is no action that can alter the trajectory of the center of gravity. There are, of course, many variables that can be controlled during the trajectory, such as the relative configuration of the parts of the body. These can make quite a difference in whether the bar falls or not, but they have no effect over the path followed by the center of gravity. One of the tricks of high-jumping is to control the body's configuration so the center of gravity passes under the bar while the body itself passes over it. That process is under continuous control all during the trajectory.

I think that competitors control what they can control: the approach, the takeoff, and the body configurations. The outcome depends on how well they are able to control those variables.

The peak height of the trajectory, perceived over dozens or hundreds of occasions, might be a controlled variable if there are things the jumper can do to affect this average peak height. The associated control system would be very slow, and would operate by adjusting many lower-order reference signals for such things as practice time, amount of effort, adjustments of form, and so forth. During any one jump, of course, this averaged perception can't be controlled. But over time, the jumper can gradually raise the reference signal for height jumped, as long as this is consistent with maintaining the necessary elements of the jump in the right forms. On the initial jump of a competition, no jumper strives for maximum height. The reference height is set comfortably above the bar, but no higher than necessary.

>I do allow the answer that the reference signal is
>discovered 'by accident', through trial and error learning.
>But that would mean that the very first time there was no
>reference that could be followed, i.e. that not _all_ behavior
>(here: peak performance) is control of perception.

I think you would have a clearer picture of the PCT approach if you kept the hierarchy in mind. The first time anything is accomplished, there can be no reference signal derived from experience of accomplishing it. At worst, one can have reference signals only for the lower-order components of perceived behavior that are to be put together in a new way. There are many possible ways for that to happen, including instruction followed by imagining the meaning of the instructions. At best, you've studied movies of someone else doing it and have some concept of the coordinations required.

On the first attempt, one seldom achieves perfect control. But the first attempt provides a perception of doing the control action, and from that experience, more realistic reference signals can be selected. Also the new control system's parameters are probably not set to the best possible values; reorganizing them takes many trials, too.

To speak of "the" reference signal being "discovered" doesn't sound right to me. A reference signal is variable; it can be set to high or low levels. In any complex behavior, reference signals must be varied during the behavior if high-level perceptions are to be controlled at their given reference levels. Even when a behavior is well-practiced, the reference signals can be set to different states within the possible range. As I said, a jumper doesn't set a reference signal for the maximum possible jump early in the competition; you don't see champion pole-vaulters clearing a 15-foot bar by 5 feet. I don't think that "maximums" have anything to do with it, anyway. The jumper simply sets a target height that is enough above the bar to clear it. When the bar is set too high, the target is still set above the bar, but now the jumper can't produce lower-level control actions sufficient to clear the bar, and fails.

If a jumper really set a reference signal for "maximum height" (say, one kilometer), there would be an enormous error signal and the output function would saturate, destroying control. To achieve maximum performance, one should set the reference signal just slightly above the level that the maximum possible efforts can achieve.

Martin Taylor (930208 16:30)--

Very nicely said on the subject of solipsism. Wrt reorganization, one of the most interesting aspects of this subject is that we can design a reorganizing system without having to know what is actually causing our perceptions, or those of the system we're designing. In fact I take that as a criterion for the proper design of a reorganizing system: it should make use only of information available inside the system, without depending on solving the epistemological problem.

Martin Taylor (930208.1700) -- (Rick 930202.1100 resp to Martin 930202 12:20)

Martin's perceptual-information experiment is, in general spirit, somewhat similar to Rick's experiment with levels of perception, in that some simple action is used to indicate whether or not something is perceived, but control of that something is not specifically being tested. It would be possible to start with either experiment and give the subject control of the variable, to make a complete control-system experiment out of it.

But I think that such experiments are useful and tell us something about perception, which is a start toward doing a more inclusive experiment in which control is investigated.

Chris Malcolm (930209) --

>So, do thermostats have intentionality in this (aboutness)
>sense? .... The argument against -- which I incline to -- says
>that this kind of intentionality is only appropriate in systems
>with a level of propostionally governed behaviour, i.e., which
>combine their symbols into collections of propositions, and
>perform some kind of reasoning with these.

If by the "aboutness" sense you mean to refer to the relationship between symbols and meanings, then I agree. As I said above, I think that the PCT approach to meaning is that words are "about" experiences in the hearer/user of the words, so they are really about perceptions, not the outside objective world. This, as I understand it, is not the standard view, which assumes that somehow we can know about both perceptions and their correlates in objective nature.

I concur with you that the use of intentionality-with-a-t in this context is unfortunate. As I said a couple of weeks ago, I think the reason is that real intentionality, that is purposeful behavior, was written off 60 years ago or so, and that people then tried to give the term a new meaning that had something to do with the original meaning without committing scientific sins. This led to great confusion and obfuscation.

>On the other hand, I definitely argue that a thermostat (by
>which I mean the whole complex of thermostat, heating system,
>room, etc.) does have a purpose, is a goal-seeking device, and
>thus is intentional (meaning, has a purpose).

Yes, absolutely. The demystification of purpose, without distorting it.

By the way, Korzybski wrote in 1933 about "intension/extension."

Avery Andrews (930210.1135) --

>A percept assumes one form inside the organism, as the firing
>rate of a neuron, for example, and another outside, as a
>complex and perhaps rather subtle property of the environment.
>To keep the Gibsonians happy, one can say that this
>transformation of form is normally achieved by means of lawful
>transformations of energy, tho I think there a cases where more
>chaotic and error-prone processes get involved.

The problem with keeping the Gibsonians happy is that you end up accepting the "lawful transformations of energy" as if we could know about them in some way independent of perception. At best, the solution you propose is simply a comparison between models: a physical model based on one kind of perception, and a neurological model based on another kind.

>One of the puzzles of semantics is how a word like `gold' or
>`plutonium' can have a meaning that is in some sense

>independent of the concepts in the brains of most of the
>individual speakers of the language. E.g. none of us could
>right now recognize plutonium & distinguish it from other
>substances without killing ourselves (though some of us could
>probably figure out some way to do it, given time & access to
>the right kind of library), but there is a sense in which we
>know what it means, & can use this knowledge effectively (to
>vote for or against making it, deciding whether or not to give
>Greenpeace money to hinder its being shipped around the world,
>etc.).

I'd rather drop the assumption that such words do in fact have any meaning that is independent of the concepts in brains of individual speakers. I think that one of the Great Illusions is that words "have" meanings. If we adopt the proposition that meaning is evoked experience, and not a reference to some objective reality, then it becomes clear that the way a person uses a word is THAT person's way and none other's. I don't think that there is any sense in which we all know what "plutonium" means, although there is a sense in which we EACH "know" what it means. We don't vote on what words actually mean, but only on the meanings that they evoke out of our own experiences. The fact that all meanings are private -- together with the common assumption that they are objective -- explains most of the woes of the world.

>I'd suggest that the meaning exists in part by virtue of
>arrangements, in the society at large, for correcting `errors'
>in the usage of the word.

Yes, I agree with this. The only addition I would make is that this process stops when all parties involved perceive no further errors -- which may be a long way from the point at which their internal meanings actually agree.

Hans Blom (930210) --

>More formal, then. Consider a car driving with constant
>speed on a narrow road with cliffs on both sides. The
>weather is a bit gusty. We consider just the position of the
>car relative to the middle of the road. Call this variable
>x. Model x as a function of time; x depends on 1) the x of a
>moment ago, 2) the way you move the steering wheel; call this
>influence u, and 3) the wind and other random influences on
>x; call these e. The model of the car's position is then
>something like (if you take a difference equation rather
>than a differential equation):

$$x (t + T) = a * x (t) + b * u (t) + e (t) \qquad (1)$$

Why not the position of the car relative to a point 1 foot to the right of the middle of the road? You're sneaking a reference condition into this argument without mentioning it.

>... where e (t) is unknown but hopefully its statistics are
>known. An often made assumption is that e=0 on average, and
>that its standard deviation is constant and known. Catastrophy
>threatens when the absolute value of x becomes too large.

In a steady crosswind, e certainly does not have an average value of zero. Assuming disturbances with an average value of zero conceals the real control problem -- such as standing up in a gravitational field.

> This imposes limits:

> | x (t) | <= xmax for all t (2)

You're converting a simple control problem into a more complicated limit problem. Why not just say that there is a position x* somewhere on the road that the driver has selected as a reference position, and that the wheel angle is proportional to (x* - x)?

>The problem is clear now: find a control law for u (t) that obeys (1) and (2).

>In linear quadratic control, the time integral of the square
>of the error is minimized. That allows an occasional large
>error, provided such large errors do not occur too often.
>Here the situation is different: even one error > xmax is
>not allowed, but otherwise you can swerve all you like.
>There is no reference signal in the strict sense, although the
>control law will show that the average position will be the
>middle of the road.

So this control law will leave the car weaving back and forth from one side of the road to the other under each gust of wind, with the driver attempting to steer only when a limit gets too close. If this is how you drive, I'm not sure I would like to be a passenger! Nor do I think that this behavior would look much like the way a real human driver steers a car.

You claim that there's no reference signal here, but you have put one in by saying arbitrarily that you will measure x relative to the center of the road, and that disturbances will average out to zero with respect to effects relative to this centered position. To explain why steering efforts are centered on this position, you have to depend on statistical averages and the assumption that disturbances average to zero. Furthermore, if you put real car dynamics into this analysis, in which the lateral position of the car goes as the second time integral of the steering wheel angle or steering effort, you would find that the car would always go off the road after a sufficient time, because it would perform a random walk away from the center position (there's nothing to return it specifically to that position).

A real human driver, however, could keep the car moving indefinitely, near to any position on the road -- the center, or the right side, or the left side -- without having to approach the cliff on either side. There's no need to suppose an unrealistic kind of disturbance to make the control system work successfully -- the wind could be gusty or have a constant superimposed velocity to either side. The only limits would be that the gusts would have to stay within a certain envelope of force and duration, so as not to push the car over the cliff before the driver could react, or push so strongly that the maximum possible steering effort could not prevent sideways motion. Of course we could allow the disturbances to exceed those limits, in which case the model car, like real cars on occasion, would go over the cliff.

I don't think your "different control law" is very practical.

Hans Blom (930211)--

>Perception is not the only human capability that we depend on
>to control our behavior. Sometimes memory will do: a child will
>stay away from a hot stove after having been 'bit' by it only once.

How could the child know its current relation to the stove without perceiving it? If it relied on memory alone, it would know where the stove WAS, but not where it IS, relative to the body.

>Sometimes 'knowledge', such as from a newspaper, will do:
>stay away from Chernobyl for a while.

Same problem. You must perceive your present relationship to Chernobyl in order to control for being "away" from it. If you're far enough from Chernobyl when you read the newspaper, there's no error and no behavior because you're already "away" from there; you need to act only when you perceive that you're too close to or even in Chernobyl.

>Maybe we have a different conception of what perception
>is. For me, perception is everything that my senses register
>and what can be derived from that. You might include memory as
>some type of 'observation' through 'inner senses'. Is that what
>you mean?

That all sounds OK to me. Perception is what we know of the world and ourselves. It exists physically as signals in a brain.

>I do not dispute that we have reference levels and that we
>use our perceptions to get us close to them. I just want to
>add something like 'negative reference levels', things to
>stay away from.

There are many reference settings that result in staying away from something. The simplest kind is a reference setting of zero. If you set your reference level for the perception of a loose tiger to zero, then any perception of a loose tiger constitutes an error, and you will act to reduce the perception of the tiger to zero by moving it away or yourself away from it.

>Freedom is a name for ranges in N-dimensional objective space
>where you can move about 'at will', because the objective
>function is flat.

You get the same result from an inverse-square function. If you keep the perception of the tiger at zero, you still have all the other degrees of freedom of movement, the only restriction (which you set yourself) being that the perception of the tiger should not depart significantly from zero.

Actually, by the way, you would probably not set the reference signal to zero, but to some small nonzero amount. If there's a tiger on the loose, you want to see a very small image of a tiger, but you definitely want to see SOME image of the tiger. It would not be wise to lose track of where it is.

Rick said:

>>The result of controlling all these variables
>>is USUALLY no accident. ... We are controlling other variables
>>IN THE HOPE that by doing so we will not be damaged.

>As a control systems designer, I must seriously object. We do
>not create control systems 'in the hope that' they function
>correctly; hope has no place in the model.

Well, you hope that somebody doesn't pull the power plug, or that the motor doesn't burn out a bearing, or that the environment doesn't become so nonlinear that your design becomes unstable, and so on. Every system, however carefully designed, has failure modes, doesn't it?

In fact, designed control systems live in an environment that's almost totally predictable, so you can be pretty sure that nothing disastrously unexpected will happen. But human beings roam free through an undisciplined environment that is far more complex than any of them can understand. That environment is also full of disturbances that can't be predicted (weather, for example) or even be sensed before they occur. Most of our "predictions" are statistical in nature; sometimes they work and sometimes they don't. So there's no way that living systems could evolve to anticipate every circumstance or act correctly every time.

There's another factor that the designer has considerable control over: the forms of the analytical functions involved in the design. Most control systems are deliberately designed with linear components for the simple reason that we can't solve the equations with nonlinear functions -- not because nature doesn't present us with nonlinear situations. In most real control problems, if you actually use the mathematical forms that fit the behavior of the environment most accurately, you find that you can't solve the equations and can't complete the design without trial and error. So we all use approximations; we fit a quadratic to the curve, instead of using a power of 2.113 which would fit better.

The human control systems have to work with the components that are given. They can't approximate.

>By 'control in the usual sense' I meant the type of
>control that is discussed mostly in CSG-L, the type that all
>Bill's models are based on. In the engineering literature, it
>is called by different names, such as linear quadratic control
>or PID control. It assumes that the plant to be controlled is
>linear (or that its non-linearities can be neglected), it
>relies on linear relations in the controller itself, and its
>objective function can be shown to be the average of the
>minimum of the square of the deviation between a setpoint
>(reference level) and an observation.

Just to correct a (justifiable) misapprehension: not all my models involve linear functions. I don't do much of the kind of mathematical analysis that real control engineers do; it's mostly working directly with simulations, and basically cutting and trying until the thing works. So I really don't put much emphasis on linearity, despite the fact that I usually use a simple linear model. In the arm model, the environment part is highly nonlinear, and so is the perceptual part that computes distance from binocular vision (the system uses the inverse of subtended angles as the distance measure). It's no harder to handle nonlinear functions than linear ones, in a simulation.

My job is actually easier than yours. I'm not trying to optimize anything -- just to match the behavior of a model with that of a real human subject. Human subjects don't exhibit optimal control, either. When there are nonlinearities deliberately introduced into the loop, the model's performance falls off -- but so does that of the real person.

Rick mentioned an experiment in which I put a cubic form into the relationship between handle position and cursor position. As the handle moved from left to right, for example, the cursor would move the same way, slower and slower, then reverse for a while, then move to the right again. The simple linear model became looser as the reversal was approached, then skipped clear past the reversed (positive feedback) segment and took up control on the other side, where feedback was negative again. The human subject did exactly the same thing. This showed that the simple model was right for the human being, too -- the human being didn't find an elegant solution for the reversal, but just skipped the positive feedback region in the same way as the model.

This sort of result encourages me to think that human control systems are not extremely complex. Sometimes they work very well, sometimes not well at all. I'm just trying to produce a model that controls as well as people do, not to produce engineering miracles.

Of course real control engineers know a lot more than I do about the design of complex control systems, and some day they will take PCT much farther than I possibly could. My job is not to compete with them or tell them their business. It's to get them to look at control in novel ways, ways that are not part of the customary approach -- and not to improve the control systems they design, but to help us understand the behavior of organisms, most of which are not control engineers, either.

>E. coli has a funny (partly random) but clever control law
>that results in what is called a biased random walk. This
>'primitive' control law serves it quite well; coli is far more
>numerous than homo sapiens.

We've modeled that one, too -- maybe Rick would send you or email you a copy of the article.

Rick says:

>>There are many other ways to model the control of x in your
>>example; the "right" way must be determined by testing the
>>model against real behavior; not against catastophe theory.
>>Most important, we don't even know that the driver is
>>controlling the variable, x, that is controlled by your model.

You say:
>We do not know that the driver wants to keep on the road
>rather than fall off the cliffs? Then what _do_ we know?

There's a difference between knowing what variable a person is actually controlling for, and realizing the consequences of controlling for that variable. A driver may be controlling to keep the lateral distance of the car from the center of the lane as small as possible. A _consequence_ of controlling for this variable is to keep the car from falling off a cliff. But there is no justification for saying that people are specifically perceiving and controlling for such consequences just because the consequences exist. One driver might be continually in a sweat about falling off a mountain road, and thus control in much the way you suggested by jerking back from contact with the edge of the road and not doing much to steer in between the limits. Another more expert driver may drive the mountain road in the same way he or she drives on any road, simply keeping the car in its proper position in its lane and not even worrying about going over the edge.

My point is that pure reason isn't going to identify the actual variable under control by a given person in a given circumstance. A guess about what someone is controlling for could be quite right, or quite wrong. The only real way to find out is to apply a disturbance to the proposed controlled variable and see whether it's resisted in the way a control system would resist it. An even better way is to match a model to the behavior and find the parameters that give the best fit, and that predict future behavior in detail. This is why we refer to the Test for the controlled variable -- because it provides a formal way of determining what is in fact being controlled, as opposed to what seems reasonable. People are not always reasonable. They don't all control for the same things in the same way. Sometimes they seem positively determined to do things the hard way. All we can do as theoreticians and experimenters is to find out what's really going on in a given person.

Martin Taylor (930211.1000) --

One nice thing about going away is that others are compelled to step into the breach. I thought your reply to Hans Blom was simply excellent.

P.S. I still have most of my hair, and I don't really pull it out any more over mere misunderstandings.

Well, that more or less catches me up to the 11th, when I downloaded the 135K of mail from my brother-in-law's house in San Diego. There's more to comment on but I'm beat and anyway others seem to be doing fine. Tomorrow I'll go through the 110K's worth that accumulated by the 16th, and then maybe things will be back to normal.

Best to all, Bill P.

Date: Wed Feb 17, 1993 7:10 pm PST
Subject: I'm a thermostat

[FROM: Dennis Delprato (930217)]

No restraints or sanity hearings needed, for....

"What's it like to be a thermostat (as opposed to being talked about being a thermostat)?"

Once, again I am reminded of the problem created when control systems are talked about from an outside perspective:

>Bill Powers (930217.1030)

>As a ... third-party observer you can say ['The signal in my
>home thermostat is about the temperature in my house']. As a
>thermostat, you could not say it. ....

How is it that despite the many depictions of heating and cooling systems, complete with thermostat, in today's literature, virtually everyone describes what's going on as control of output? I suggest someone prepare a kindly little essay (?) spelling out how easy it is to be deceived when one looks from the outside in and even does a bang-up job of describing what they observe. Point out how one gets a very different picture when one "takes the viewpoint of the thermostat." Seems like PCT is getting closer and closer to Stephenson's Q-methodological thinking that stresses how psychology ignores the person's point of view, instead imposing the observer's point of view on the person and calling it understanding the person.

Bill reminded me that even the thermostat "has a point of view," and this is what he (Bill) is concerned with.

Date: Wed Feb 17, 1993 8:13 pm PST
Subject: deafferentation paper (good guys)

[Avery Andrews 930218.1501]

A useful paper if you expect to get into arguments about deafferentation:

Sanes, J. (1990) `Motor Representations in Deafferented Humans: A Mechanism for Disordered Movement Performance.', in _Attention & Performance XIII_, 714-735. The gist is that partially deafferented people are klutzy. Put more scholarly, people with large fiber sensory neuropathy show various deficits that suggest that they have an impaired effort sense, especially for small amounts of effort, which shows that there's more to the muscular effort sense than efferent copy. Sanes also argues that the Bizzi equilibrium point hypothesis is just wrong for human limb positioning (he doesn't challenge the monkey pointing stuff, however, but then the circumstances under which those gestures are produced are very special).

One for the goodguy list, I'd say. Avery
Date: Thu Feb 18, 1993 8:42 am PST
Subject: Re: re: re: re: control article

[Martin Taylor 930218 10:40] Hans Blom, 930217

>Engineers and psychologists are not close neighbors. They speak different
>languages, have a different culture and work on different problems,
>although it is fascinating to discover similarities.

Well, they are close enough neighbours in this head. My certificates say I am a certified Professional Engineer in the Province of Ontario, and I have a Ph.D. in experimental psychology. Close enough? I always thought psychology was essentially a problem in engineering, which is why I seemed to switch fields (according to society--I never thought I switched). I had thought of doing my Master's in control engineering, never realizing its connection with psychology at the time (that realization came with reading CSG-L). When I had newly graduated, it seemed that many of the prominent psychologists had an engineering training, so I am in no way unusual.

\>>If you can predict exactly what the results of actions will be, you don't
\>>need control. If you can't predict at all how your actions will affect
\>>the world, you can't perform control.
\>
\>Control engineers have a broader conception of control than you seem to do.
\>Control does not necessarily imply feedback.

Fine. There seem to be two language conventions here. Within CSG-L, the word "control" has always been used to distinguish situations with feedback from those without. In your community, that distinction is not made. Use of the same term with strongly but subtly different meanings is a surefire recipe for confusion, and I suspect that some of the come-back to you has been based on that confusion. There's no right or wrong here, but there is a question of communicative effectiveness. Either you must convince CSG-L readers and writers to use the more extended meaning, or you must recognize that within CSG-L postings the more restricted meaning is probably intended.

Where there is no feedback, CSG-L tends to use terms such as "affect," "influence," "linkage," and the like. The car's steering wheel is linked to direction of the rubber-tired wheels. There is some degree of control (CSG-L-sense) here, in that the reactive forces could be construed as representing feedback, but I think that interpretation would really be pushing it beyond any reasonable extreme. If, however, the steering wheel angle is converted through an optical sensor and a transducer, which provides the force to move the road wheels, even that feedback is lost, and you get a real outflow "control" (Control-engineers' sense).

\>Regrettably, non-feedback control is possible only if the system to be
\>controlled in invariable and not significantly subject to disturbances.

The presumption on which PCT is based is that these circumstances are rare when we are dealing with living systems. I add my personal claim that this is so because human chemical systems are thermodynamically unstable at the temperatures and in the energy flows within which they exist. However, given that there is low-level control, then sometimes higher-level systems can with high probability get results near the intended results without feedback. In this, I think I have a disagreement that is more apparent than real with some others on CSG-L.

\>>Avoidance is that the ECS brings its percept to a value different from
\>>the reference in either direction.
\>
\>Can I equate "a value different from the reference" with "a new reference"?
\>If so, I would not call it avoidance.

No you can't. I went on to describe an avoiding system, and I don't think you could have asked this question if you had read it. An avoiding system has a high gain POSITIVE feedback loop for small errors, the gain decreasing to zero for large errors (the term error here should be construed in its PCT sense, as a deviation between perception and reference, with no connotation of "mistake").

>An adaptive control system can learn the relations between its
>actions and their effects by generating an action and observing its effect.
>If a non-zero effect results, control becomes possible.

Sure. That's the point about having some predictability. All your statements later in your posting about ways of learning are valid, and have been discussed here. More about them below.

>>(1) If two actions are incompatible in the world, the control loops that
>>involve them are mutually inhibitory. Only one of the relevant perceptions
>>(or neither) will actually be controlled. We have had some discussion of
>>this in the past. I like to look on it as a question of loop impedance,
>>in which the coupled real-world systems affect each other's impedance and
>>hence the gain of the loops in which they participate.
>
>The term "impedance" assumes a linear and static world. What if the world
>kicks back, as in hysteresis?

Not always. Even as an undergraduate I learned about regions of negative and positive impedance in tube and transistor circuits. That's how you design oscillators. There's no problem there with hysteresis. Have you tried magnetic amplifiers? It's a bit restrictive to limit your definition of "impedance" to either static or linear.

>>(2) The second aspect of decision is in the distribution of output from an
>>ECS. It is conceptually possible for an output function to alter the
>>distribution of its outputs as a function of the error. That could be seen
>>as decision (I am far in space and near in time to somewhere I want to
>>perceive myself, so I take the car, whereas if I were near in space and
>>far in time, I would walk).
>
>I would call this a selection between different possibilities of action.

That's exactly what I would call "decision." Do we have another source of confusion based solely on a different dictionary? What do you mean by "decision?"

>Luckily, humans are wired in such a way that they can sense their outputs;
>this is called the "body image".

No. They sense inputs from many sensors, some of which are detecting the conditions inside the body. But it is quite possible (and discussed in BCP) that the anticipated (imagined) effects of outputs can be used as inputs through what you call models.

You talk about two kinds of learning in adaptive control theory. Within PCT there are at least 12, but your types of learning are not related to those 12. What you are talking about is ways of acquiring information that permits learning, not about learning itself. By the way, have you noted the relation between your two types and the distinction between fixed and variable effects in ANOVA? Active learning permits generalization, passive does not, and this is reflected in the perceptual consequences. Active exploration is more likely to lead to a perception of "an outer world," as Jim Gibson well knew. The ANOVA analogy helps one to understand why.

>I think that by now I understand what PCT is about. I have followed and
>enjoyed the discussions for more than a year now, mostly quietly. Once in a
>while I grab the chance to vent some of my ideas, which are more or less
>related, hoping for a useful reply -- usually not in vain. Reconciliation
>is not what I look for; I find that friction -- clashing points of view --
>generates much more creative energy.

I find, as you must, that when I think I understand PCT, and make some didactic comment, there are always people with a better understanding who are willing to show me the error of my ways. This is all to the good. But friction generates more heat than light, and is not a good way to develop higher levels of organization.

Martin

Date: Thu Feb 18, 1993 8:52 am PST
Subject: Re: control article, focus

[Martin Taylor 930218 11:20] Rick Marken 930217.1100

>>The slide example is not relevant, except as an instance in which partial
>> prediction is possible.
>
>I think the slide is an excellent example of control with zero pre-
>dictability of the effect of the control action; another easy to
>demonstrate example is tuning a radio. There is no way to predict
>(better then 50%) whether a clockwise or counterclockwise turn of
>the tuning dial will improve or degrade the degree of tune -- you
>just have to turn and see what happens.

There's a misunderstanding somewhere. Once you have turned the knob one way or the other, you know that your perception is getting closer to or further from the reference value, and you can predict perfectly (well, almost) that if you keep going the same way it will continue to get better (or worse). I'd call this perfect prediction except for a one-bit uncertainty. It is not relevant to the situation of no predictability in which control is impossible.

>Focusing and tuning just show that there are certain kinds of control that
>CANNOT be done by a single level control system.

Yes.

>We can tune radios because we can perceive and hence) control whether
>we are "approaching" or "moving away" from tuned (focused).

And can predict that we will continue to do so if we keep producing the same kind of output (unless we pass the reference value).

I'm trying to write the "Information leads to PCT" paper in odd spare moments. Misundertandings such as those of Gary and Rick illustrate why it is a difficult thing to do. In writing it, I have to explain concepts that in my earlier postings I had taken for granted would be understood. Now here's another of the same kind that will have to be incorporated in that paper. All to the good, if we can bring these potential problems into the open before they cause real communication difficulties.

Martin

Date: Thu Feb 18, 1993 10:22 am PST
Subject: Re: Catching up on many subjects

[From Oded Maler 930218.1900-ET] From Bill Powers (930217.1030)

* The Little Man can already do this. By sending a sine and cosine
* wave to the reference inputs of the x and y visual control
* systems, the program can make the fingertip draw a circle
* continuously around the target, whether the target is stationary
* or moving. By specifying more complex waveforms one could produce
* any pattern of movements relative to the target. I don't know how
* to model a perceptual system that can perceive the patterns of
* writing, so I can't build the level of control that produces
* those waveforms in the manner of a control system. But the model
* provides for reference signal inputs that will make the fingertip
* movements follow any prescribed pattern.
*
* There isn't any "noise between levels" in this system. Why should
* such noise exist?

If you stay with circles, it might be interesting to have a "motor program" (your version of the concept, that is with continuos perceptual feed-back, no access to "objective" variables) that draws circles robustly, which means

1) different centers (you say you already have), different radius as a parameter.

2) higher-levels disturbances, constraints, e.g., asymmetric change in the mechanical parameters of a joint, or some constraints on the "workspace" of the joints (not much can be done in this direction with the current no. of degrees of freedom without restricting the range of possible end-point behaviors). Add "friction" to the fingertip according to the direction of its movements, etc. If you manage to write a "program" that succeeds in making those circles (and later, "A"s) under different kinds of circumstances, and still they will look to a observer as good approximation of circles or "A"s, you succeed in demonstrating the PCT notion of a motor program.

Oded Maler

Date: Thu Feb 18, 1993 10:29 am PST
Subject: meaning, control, prediction

[From Rick Marken (930218.0900)] Bill Powers (930217.1030) --

>I'd rather drop the assumption that such words do in fact have
>any meaning that is independent of the concepts in brains of
>individual speakers. I think that one of the Great Illusions is
>that words "have" meanings. If we adopt the proposition that
>meaning is evoked experience, and not a reference to some
>objective reality, then it becomes clear that the way a person
>uses a word is THAT person's way and none other's. I don't think
>that there is any sense in which we all know what "plutonium"
>means, although there is a sense in which we EACH "know" what it
>means. We don't vote on what words actually mean, but only on the
>meanings that they evoke out of our own experiences. The fact
>that all meanings are private -- together with the common
>assumption that they are objective -- explains most of the woes of the world.

Why isn't this approach to meaning more obviously a part of cognitive psychology? Or is it, and I've just never noticed because it was never expressed in a way that evoked the kind of imagery in me that Bill's description above did? Law, religion, and other human foibles are based on the idea that words have meaning "in themselves"; why else would lawyers constantly argue about what sentences (laws) "really" mean; why would religious idiots try to explain what sentences (god's word) "really" mean. If people could be disabused of this transparently idiotic idea (that words have meaning) and be educated to understand that words only "mean" the experiences that they evoke in each individual -- experiences that are likely to be considerably different across individuals -- maybe people could start approaching the problem of communication from a whole new perspective; one aimed at convergence on common experience rather than common wording.

Bill to Hans Blom:

>Of course real control engineers know a lot more than I do about
>the design of complex control systems, and some day they will
>take PCT much farther than I possibly could.

No offense, Bill, but this sounds somewhat disingenuous. You ARE a real control engineer and I think it must be obvious to anyone reading this list that your knowledge of control is equal to (in my opinion, it is orders of magnitude greater than) that of so-called "real" control engineers. Just because people have a degree or a certificate in field "X" (so that they are a "real Xologist") does not mean that they have a particularly deep understanding of X (just try replacing X with "psychologist").

Martin Taylor (930218 11:20) --

>There's a misunderstanding somewhere. Once you have turned the knob
>one way or the other, you know that your perception is getting closer
>to or further from the reference value, and you can predict perfectly
>(well, almost) that if you keep going the same way it will continue
>to get better (or worse).

I thought that someone had claimed that control is only possible when the effect of the output on the input is perfectly predictable. The "tuning" example shows that control is possible if the effect of output on input in not predictable IF you have a higher level system that can deal with the result of the unpredictability. I don't see how the predictability of the effect of the tuning knob AFTER you start turning makes up for the fact that when you START turning the knob the effect is perfectly UNPREDICTABLE. If the direction of turn you select makes the tuning worse, it will PREDICTABLY continue to make it worse -- so there is predictability and NO control. My point is that "predictability" of the effect of output on input is not a necessary requirement for control; there are ways to build systems that deal with this unpredictability. A better example, perhaps, is e. coli. Here, the effect of an output (the direction of movement after a "tumble") on the input (sensed gradient of attractant) is COMPLETELY unpredictable -- because the direction is random. Moreover, even during movement it is impossible to "predict" changes in the gradient because these changes depend on the angle of movement relative to the center of the gradient (to say nothing of the effects of disturb- ances, such currents in the medium, that could alter the direction of movement while the bacterium is "swimming"). As long as the system can keep changing outputs based on the degree to which perception deviates from intended perception (changing when the deviation is large or increasing and holding pat when the deviation is small or decreasing) it doesn't really need to be able to "predict" the effect of output on input.

Just another wonderful benefit of the fact that control is the control of PERCEPTION, not OUTPUT.

Best Rick

Date: Thu Feb 18, 1993 12:23 pm PST
Subject: More catching up

[From Bill Powers (930218.0730)] Oded Maler (930211.1100 ET) --

Excellent plan for developing the paper, including the 6-month deadline. Some suggestions:

>1) Definition of the domain - low-level sensory-motor behavior
>(up to the level of pointing-like behavior).

I think this should be preceded by a more general statement about the scope of PCT, lest we perpetuate the popular idea that control theory applies ONLY to low-level sensory-motor behavior. Here's a fragment for the ultimate editor of this paper to chop up:

The theory of closed-loop control applies to behavior at many levels of organization. It is needed in all circumstances in which consequences of behavior remain stable or repeat while ongoing changes in environmental conditions require that different actions be employed to produce or maintain the same result. In a typical control-system model of behavior, part of the model is an independent environmental disturbance that affects the same outcome that the actions of the behaving system affect. Such a model senses the state of the outcome, compares the resulting perceptual representation with a reference criterion, and converts the difference into signals that vary the output forces. Hence we refer to this application of control theory as "perceptual control theory" or PCT. In such a model it is a perception of the outcome, not the action, that is under control by the organism.

The result is that the net outcome, the summed effects of output forces and independent disturbances, comes to a specific condition and is maintained in that condition. The result is not "homeostasis" (Cannon), but "rheostasis" (Myrsovski sp??), for the reference criterion itself can be adjusted by higher levels of control. This is the basis for a _hierarchical_ theory of control, or HPCT.

In this paper we restrict our discussion to lower levels of perceptual control and their relationship to other theories of organized and coordinated motor behavior....

>4) The difference between control-engineering/robotics criteria
>and PCT criteria for the adequacy of models.

I think we need to distinguish clearly between closed-loop and open-loop control, as engineers seem to include open-loop systems in their definitions of control. Something like this, leading into a discussion of a PCT model for arm movement:

There are two ways to design automatic systems to stabilize external variables against disturbances. One of them is closed-loop control, which we use here. The other can be called "compensatory" control.

In a compensatory control system, the basic system is designed to convert a command input into an output that has the desired effect on a variable in the environment. For any simple system of this kind to work, the affected variable must be protected against disturbances arising independently of the system's own output forces. When such protection is impossible, the sources of the disturbing effects must be sensed, and the sensory signals must be used as the basis for calculating compensating changes in the output of the system. In this way the direct effects of the disturbances on the environmental variable are cancelled by superimposed changes in the output of the active system. The output of the system thus consists of a basic output calculated to produce the desired effect in the absence of disturbances, plus a variable component that is a function of any independent disturbances that may arise. This was the concept of control that W. Ross Ashby finally accepted in the 1950s, and that is the basis for many modern approaches to the analysis of human motor control.

There are two serious drawbacks to the compensatory model. One is that not all sources of disturbance can be sensed before their effects on the final outcome begin to occur, yet in real behavior the effects of such disturbances are resisted. The other is a precision problem: to compute the outputs required to achieve many kinds of physical effects, one or two time integrations are typically needed, and such integrations are hypersensitive to computation errors and changes in initial conditions. The latter problem appears even when the sources of disturbance can be sensed. The sensing itself must be quantitative to the highest degree, and both the computed compensations and the response of the physical actuators to the modified commands must be extremely precise.

In most kinds of ordinary human behavior, the accuracy of sensing and the calibration of motor output equipment is only moderate -- the accuracy and repeatability are a few percent at best. We must ask whether they are good enough to justify a compensatory model of motor control.

Consider the problem of moving a fingertip rapidly from one position to another by a simple shoulder joint-angle change of 60 degrees. The change in angular position is the second integral of angular acceleration (the immediate effect of muscle tension), so that p = 1/2at^2, where p is angular position of the half-way point, a is angular acceleration, and t is the elapsed time. In order for the final position to be within one degree of the desired 60 degree movement, the product at^2 must be calculated and produced with an accuracy of 1.6 percent during acceleration and deceleration phases combined.

This is for a single movement. If we now try to generate a series of 10 consecutive movements between fixed angular positions at the same speed, the elapsed time will be 10 times as long. The effect of time errors will go up by a factor of about 20 and the effect of (random) acceleration errors will increase by roughly a factor of 3. Now the precision requirement is something like 3 parts in ten thousand, patently impossible for a system made of nerve and muscle.

This situation is somewhat improved if we use a realistic model of the muscle. A command signal alters the resting lengths of muscles; with opposing muscle pairs, the command signal specifies an equilibrium position (Bizzi, Kelso, etc.). If there is sufficient damping in the muscle, and if the series spring component of the muscle is sufficiently stiff, the accuracy of any movement will remain as constant over time as the muscle properties.

However, the required spring constants and damping are not present in muscles. Holding an arm out horizontally requires a certain force to counteract gravity. Moving the arm through a 60- degree arc in 0.2 seconds (which is quite possible), however, entails muscle forces 10 to 20 times as large as the "resting" forces, and moreover is known to involve strong deceleration contractions of the muscles opposing the movement. If muscle damping and spring constants were adequate to explain this kind of movement, there would be no deceleration contractions.

If passive muscle properties can't account for behavior during fast movements, a compensatory model must introduce inverse dynamic calculations so that driving waveforms can be adjusted to provide the required decelerations and apparent damping. We are then back to the problem of hypersensitivity to initial conditions, for such calculations require many time integrations.

There is clearly reason for doubt that a compensatory model can adequately explain simple motor acts. We are therefore led to consider the other control model, the one based on closed-loop feedback control.

Bob Clark (930212.1030 EST) --

>the reason to go to Paris is omitted.

I take it for granted that some higher system is specifying the reference condition of going to Paris.

>I'd like to report that my life "while getting it" [PhD]
>included a minimal percentage of "hell."

Glad to hear it. I think that experience would be more common in a science like physics.

>The new "aiming point" is the "new target" for the gun crew. The target
>for the crew is no more no less than that ordered by the commander.

Yes, there are two levels of control involved here. Considering only the commander's level, the target always remains the same: the position where the shell is intended to land. The error is the amount by which the gun-crew misses the target. The commander must alter the target position given to the gun crew SLOWLY, however, to avoid treating dispersion in the pattern of shots as a systematic error.

>HIGH LOOP GAIN
>It seems to me you are following events around the loop,
>resembling open loop analysis.

This is indeed difficult to convey accurately. Loop gain is in fact the product of all amplification factors encountered in one trip around the closed loop, so calculating it seems like following events around the loop. To get high loop gain when there are transport delays in the loop, one must also use dynamic slowing of error corrections, a low-pass filter. With the filter in place, the behavior of the system at low frequencies is just as though no transport lag existed. So even though all real system do entail such lags, they can be neglected! A difficult point to get across.

>An over-correction might occur if the gun controls were not
>properly calibrated. As I understand it, a bracketing
>procedure is often used to calibrate the gun controls.

But the point of closed-loop control is that the output function does NOT have to be accurately calibrated; it can even change its calibration during the control process. Of course it mustn't change signficantly over the space of two or three shots.

>"Whether open loop or closed loop analysis is appropriate
>depends on the time scale selected."

I don't think this is quite right. If one does an analysis on a short time-scale where delays are visible, but neglects dynamic effects, a control system with a loop gain more than -1 will be incorrectly predicted to be unstable. The existence of large negative loop gains can be explained in a sequential analysis only if the proper low-pass filtering is taken into account -- and it is usually not taken into account in open-loop analyses.

Consider a control system in which the controlled quantity is equal to the output of the system, the input function has a gain of 1, and the output function has a gain of 100. If there are lags in this system, as there are in all real systems, you would predict on that basis alone that the system would go into violent overshoots increasing without limit by a factor of 100 on every iteration. But now add a slowing factor that follows the rule "on each iteration, calculate the new output, and then let the actual output change by 1% of distance from the previous amount to the new calculated amount." This is a low-pass filter that does not alter the final steady state. The system will suddenly become stable; in fact, it will bring the error down to 1% in a single iteration! The effective long-term loop gain is still 100, so errors will be kept small over the long run.

If you try to eliminate the overshoots in this sequential system by just lowering the output gain to less than 1, the result will be stability, but the error remaining at equilibrium will be 50% of the value of the reference signal on the average. So you get stability, but almost no control. The high-gain system with the low-pass filter will counteract errors slightly more slowly, but will eliminate 99% of their effects. The low-gain system without filtering will counteract disturbances instantly, but will cancel only half of their long-term effect.

So there is a difference between closed-loop and open-loop analysis that is independent of the time-scale.

>Quantitative vs Qualitative: Both the open loop and the closed
>loop descriptions can be described in either qualitative or
>quantitative terms. The differences are in the viewpoint, and
>the purpose for the description.

If the purpose of the description is prediction, then the qualitative analysis always loses out: its predictive power is very low.

>Your example of the fire hose for "Lagged" control seems to
>work very well. But I don't think the fire chief cares which
>form of control it is as long as the water lands where HE
>specified. The chief uses a time scale of, perhaps, minutes vs
>the seconds needed for the water to flow.

My point was that all components of a closed-loop system of this sort are operating literally simultaneously; they don't take turns acting, with no action between. This is how the nervous system works; sensors are generating signals at literally the same time that actuators are producing forces.

>As suggested, an "exotic kind of conflict" occurs when the time
>scales overlap. If the spotter is repeatedly moved to a new
>position before the operations from the preceding position has
>been completed, a loss of accuracy (perhaps temporary) results.

With proper design, the system would work better if the spotter were moved immediately rather than waiting for the previous results to come in. This would be the right strategy if the calculations were being continuously averaged over several shots, as would be necessary to distinguish random from systematic errors.

Avery Andrews (930212.0906) -- cribbed from Taylor

>On a somewhat different note, I wonder if `feedback' isn't
>actually a dirty word (two four-letter ones). The problem with
>it is that it suggests that the organism's perceptual functions
>are able to draw a distinction between the effects of what the
>organism does, versus the effects of the disturbances, ...

It's easier to understand if you deal first with the disturbance- free case, so the input is strictly a function of the output. Introducing disturbances then shows that the system doesn't need to make this distinction; it just acts on the controlled variable to restore it to the reference condition no matter why it deviated from it.

Martin Taylor (930212.1315) --

>In the imagination loop, then, the results of actions and of
>disturbances can be independently perceived.

This is a nice point, one that Wayne Hershberger made some time ago on the net. There's another way to derive this phenomenon without invoking imagination. When you lift a suitcase of unknown weight, your position-control systems automatically create the amount of force needed to bring the suitcase to the new reference position. The signals from the muscles that measure force for the spinal control loops are also available to higher systems. So the higher systems can use these signals as a perception of the "weight" of the suitcase. So what we label "weight" is really "effort."

Your proposal is another argument for the modified model we've been kicking around for a year or so, or maybe another related version of it. Can you draw if diagram of just how the various signals would be treated?

Mark Olson (930212.1433) --

The "place" cells are an interesting idea. It seems to me that they would belong in the parts of the brain where the spatial arrangement of peripheral receptors is mapped into a region of the brain, as repeatedly occurs in the visual systems. This is quite a different picture from the standard PCT model in which position would be represented by a single signal of variable magnitude. It's probably more correct, too, although the final effect would be the same.

The biggest conceptual problem with place cells is how place is then represented perceptually. Soemthing has to know that it's THIS cell and not THAT cell that's being excited; one becomes cautious about infinite regress. Your proposal is more in line with Martin Taylor's concepts of distributed perceptions than with my proposals. This is hard stuff to think about.

>When I become "conditioned" to blink to the tone without the
>puff of air, how did that happen.

See Wayne Hershberger in _Volitional Action_ for a discussion of PCT and classical conditioning. Also, maybe Dick Robertson would favor the net with a summary of his experiences with the eye-blink reflex and its conditioning. Lots of things reported as solid facts in the literature just ain't so.

Bob Clark (930215.1155 EST) --

>what about those who are not as knowledgeable? How do they
>manage? What are the categories, etc that they form and live by?

I see your point and agree that it has to be considered. My levels are intended to describe categories of experience that all people (and even animals) employ without any training or knowledge. All people perceive and control relationships, by my account. They also perceive and control categories, sequences, logical functions, etc., not by thinking about it but simply by

having the world presented to them in such terms by the basic equipment of their own brains. I don't know how to put it better than that.

>Bill, I am troubled by your move from your Fifth Order, Control
>of Sequence, to discussion of "concepts." Are these concepts
>derived from combinations of lower order perceptual variables?
>If so, how? And which?

The levels as of now (9302128.1032) are (1) intensity, (2) sensation, (3) configuration, (4) transition, (5) event, (6) relationship, (7) category, (8) sequence, (9) program, (10) principle, and (11) system concept. Each one, when analyzed into components that are not just smaller groups of the same level, proves to be a function of perceptions of the next lower level (or lower still). So a system concept like physics is drawn from perceptions of many physical principles, while principles are drawn from perceptions of many specific logical/mathematical operations, and so on down the list.

As to HOW a perceptual function of one level combines lower-level perceptions, I have no idea. The nature of the functions must be very complex at the higher levels, or at least of a kind that we can't analyze now. The apparent dependencies were arrived at from analysis of experience, much as we can see that configurations are composed of sets of sensations. Also it was helpful to ask how we would go about maintaining a perception of any given level against disturbances -- how, for example, we would maintain the principle of honesty. To perceive ourselves as honest, we set reference signals for certain programs of action and thought which we call reasoning or analysis or procedures. None of this is very firm; I'm just reporting how it seems to me after as close an inspection as I can carry out. Other people's opinions are obviously needed.

I chose the term "system concept" with the emphasis on "system," not "concept". In my view, "concept" falls within the range of meaning of "perception" because it's something we can experience as occurring or existing. I could have said "system perception." It just means the sense of an organized entity of some sort being present, the kind that is composed of principles, generalizations, heuristics, characteristics, whatever you want to call them. Perceiving a specific person whom you know well leads to this sort of system concept or perception -- the impression of a particular person, a personality, a system. Shoot, how am I suppose to be more specific about an idea that's not very clear to begin with?

>To me, a physicist, it is not "a" concept, rather it is a
>specialized language, including its own special words, syntax,
>etc. It is an assemblage of definitions, observations,
>methods, procedures, formulas, derivations, etc etc.

Yes, that's what I mean by a system concept. The very fact that you can, without enumerating, refer to all its components as some sort of bringing-together into an "assemblage" of a variety of more specific elements shows that you have formed a conception of physics as a unified system of ideas, definitions, observations, methods, procedures, etc., with the "etc." indicating that the picture includes much that is not enumerated. "Physics" is clearly a system concept quite different from "religion" or "family." Enumerating the lower-level details of these other system concepts would entail quite a different list.

When you say "I am a physicist," the "I" being indicated is associated with the system concept of physics. For the moment, the center of awareness is operating from that position. But when you say "I am a father" the system concept is the one we refer to as "family," and the "I" now takes on new characteristics associated with a different system concept.

Or at least that makes a good story.

As to other differences, let's just go along with them for now. I'm feeling a bit overloaded.

Oded Maler (930216.1330) --

>The article contains, in particular, description of experiment
>in arm trajectory modification, that is, the subject is told to
>point to some target, and while moving (not seeing his own hand
>to prevent visual feed-back) he is told to point to another
>target. One suggested explanation (using terminology of
>"plans", but that's not the point) that the resutling
>trajectory is not a result of "aborting" the first one and
>starting the new one, but that it is a result of "vector
>addition" of the initial trajectory with a time-shifted
>trajectory from the origin to the new target. It might be
>interesting to see how it such things are explained based on
>Little Man's phisiology,

This would be equivalent to setting a kinesthetic reference signal for fingertip position (second kinesthetic level in Little Man Version 2), and then switching it to a new value in the middle of the movement. The result would be (is) a finger trajectory that starts off in one direction, then veers to the new target position. No plan necessary.

Hans Blom (930216) --

You're a pretty tolerant fellow, for an engineer. I'm glad that you're putting some honest effort into grasping the PCT point of view. We really need engineers, better ones than me, in this effort.

>I know by now what you mean by the mantra 'organisms control perception'.
>As so often with jargon, it is an abbreviation for a whole philosophy and
>only understandable for those who have gotten to know that philosophy. It
>is right, from a certain perspective. From another perspective, organisms
>control their outputs.

As several others have said, this isn't really jargon or "in" talk, but it is a problem with word usage.

When I think of the "output" of a system, I mean the physical effect on the environment that is due to the actions of the behaving system ALONE. In the human system, this would mean muscle tensions, because that's that last place in the chain of outgoing effects where environmental disturbances can't get into the process and alter the consequences. Measuring the consequences any farther from the nervous system can give a false impression of what the nervous system is actually doing.

In a servo system, with this understanding of "output," I would not call the output of a motor the shaft position or speed, but the torque applied to the armature of the motor (at low speeds, anyway). Only that torque can be varied by the active system without regard to what the environment is doing. Only the torque output gives an accurate indication of the electrical output of the control system. The shaft position or speed will depend on the torque AND on external loads and disturbances, so can't be used to indicate the output activities of the control system by itself (especially if the loads and disturbances aren't predictable).

So this is more a matter of labeling than ideology. I'm sure you would agree that a servomechanism doesn't control the torque applied to the armature of its motor, but only some consequence of that torque measured farther downstream in the causal chain. As disturbances come and go, the servo system VARIES its output torque, but doesn't try to maintain any particular torque (unless torque itself is being sensed and controlled, which isn't the most common case). The torque has to be free to vary if disturbances of position or speed are to be counteracted.

The "control of perception" part is also a matter of labelling. I think you'll agree that in order to control an effect of a system's actuator output (to distinguish it from "outputs" farther along the chain), that effect must be monitored by a sensor and accurately represented as a signal. The more accurate the representation, the more accurate the control can be. Furthermore, if the sensor characteristics change, the SIGNAL will still be brought to a match with the reference signal, but the VARIABLE it represents will no longer be maintained in the same condition. If the temperature-sensing element of a thermostat goes out of calibration, the thermostat will still think it is controlling the same temperature, and will keep its movable contact nearly at the same position as before, but the room temperature will be controlled at a different level.

The only aspect of a control loop that is under reliable control, therefore, is the sensor signal. The external counterpart of that signal remains under reliable control only as long as the sensor keeps its calibration accurately. So if we had to pin down any one aspect of the loop to be "the" controlled aspect of the situation, we would have to choose the sensor signal. Sensor signal = perceptual signal; hence, control of perception.

I think that my way of defining output and control is the least ambiguous. After all, if you define output at a place where disturbances can have an effect, you can't reason backward to the power or force output of the control system just from knowing the state of the variable called "output" because disturbances are contributing an unknown amount to the state of that variable. It seems strange to me to define output in such a way that by knowing the output you can't deduce what the control system is putting out. I don't object to looser usages for the sake of convenience, but when we want to avoid misunderstandings, I think my usage is the least ambiguous.

>What is the control organization of a virus over and beyond that of a
>hydrogen molecule?

The virus (apparently) senses conditions in the host and acts on its relationship to those conditions by altering its own conformations. I don't know if any true control processes are involved, but the possibility is clearly there.

Even DNA, which viruses contain, produces enzymes that restore disturbances of the coding to a preferred form. I don't know if virus DNA contains these "repair enzymes" as they are called, but if they do, that is certainly beyond the capacities of a hydrogen molecule.

>Perception is controlled by actions; actions are controlled by perception.
>Remember the loop!

Let's not confuse "control" with "affect." Control entails bringing a variable to a specified state and keeping it there. Perceptions don't bring actions to a specified state and keep them there. It's the variations in the actions that bring perceptions to specified states, despite disturbances that tend to change their states. If you add a disturbance to the actuator output of a control system, the control system will alter its own output effects, not keep them the same.

In ordinary environments, the loop is assymetrical. There is power gain going through the organism, power loss going through the environment. The part of the loop with the power gain does all the controlling.

Please, show me the diagram of this model; the only sense I can make of it is that references can be set based on remembered perceptions.

The idea is that reference signals ARE played-back recordings of perceptual signals, in organisms. This will remain only an idea until someone does the implied experiments, to see if reference signals are ever set to values that have never been experienced. I make no predictions one way or the other.

There is no reference for distance, no single 'far'. Control systems work with _numerical_ values. Therefore 'at least as far as X' is required, with the value of X depending on other perceptions.

If I ask you to hold your finger two feet from your eyes, I believe you could do that (according to your perception of "two feet"). As you say, control systems work with numerical -- quantitative -- values.

The word "far" is just a word. Before it can lead to any control process it has to be converted to a quantitative value in the dimension of distance, against which a perception of distance can be compared. In fact, I suspect that the conversion involves the reciprocal of distance, because it's hard to imagine a reference signal large enough to specify "infinitely far", and there's a simpler interpretation. In my "crowd" model I used the reciprocal of distance (inverse-square, actually) and called it "proximity." Now it was easy to get avoidance behavior just by setting the reference proximity to zero, an easily attainable value. And "nearness" became a nonzero value of proximity, with some largish but reasonable value implying a very close approach (limited by the size of the "people").

We can't know a priori how any given person translates words into specific perceptions and reference signals, but by using the Test we can refine our guesses. Chances are that the mishmash of words translates into much simpler and quite quantitative reference conditions.

>My reference for pain is zero. Having a zero perception of pain, however,
>does not tell me how far away I am from pain.

"Pain" is not an either-or sensation; it begins at zero and rises from there, with some level being considered "too much" and calling for action to reduce it. Most "pain", I suspect, is really just an ordinary sensation, like the sensation of having a fold of skin squeezed. We use the word "painful" to describe the sensation when it exceeds a certain level (although for some sensations, like stimulation of a nerve in a tooth, that level seems to be set to zero). An example in point is temperature. Moving your finger close to a hot soldering iron results in a nonzero temperature sensation, which you can control relative to any level of desired warmth by moving your finger. This sensation exceeds the tolerable limit before you actually touch the hot metal. I'd model this as two control systems, one fine and one coarse. The coarse control system has a large dead zone but high gain; the fine one has no dead zone and moderate gain. That would seem to reproduce the phenomenon. Is that the way it really works? Who knows? But there doesn't seem to be a conceptual problem here.

>If the patient is in control, drug delivery proceeds beautifully.

My wife Mary has some experience with this, and confirms your conclusion. Much less drug use, much less pain. And a whole lot less fury at the medical system.

Best to all, Bill P.

Date: Thu Feb 18, 1993 12:50 pm PST
Subject: Re: Mantra Schmantra

[from Dick Robertson] (930218.1400)

Nicely put Rick. I would make one small quibble. If we correctly named the system, Temperature Control System, we have something more comparable to analogize with the control systems we are interested in in living beings. In this view the thermostat is a compound component of the feedback function and the comparator function and the furnace is the output mechanism. The continual re-] though. And that is, how does the output function know how much output to put out in order get the controlled perception to its refereeence state? Examining the armdemo it seems to me that the output always simply hunts. Is that right? Is there more to it than that?
All the best. Dick Robertson

Date: Thu Feb 18, 1993 12:51 pm PST
Subject: Re: control article

[From Dick Robertson] 930218.1410
To Martin Taylor - I enjoyed your reply to the control article. I wonder if you would have somelthing to say to the question about output that I posted in reply to Rick's reply? Thanks.
Best, Dick Robertson

Date: Thu Feb 18, 1993 3:55 pm PST

Subject: Re: More catching up

[Martin Taylor 930218 18:30] Bill Powers 930218.0730

>Martin Taylor (930212.1315) --

>>In the imagination loop, then, the results of actions and of
>>disturbances can be independently perceived.

>This is a nice point, one that Wayne Hershberger made some time
>ago on the net. There's another way to derive this phenomenon
>without invoking imagination. When you lift a suitcase of unknown
>weight, your position-control systems automatically create the
>amount of force needed to bring the suitcase to the new reference
>position. The signals from the muscles that measure force for the
>spinal control loops are also available to higher systems. So the
>higher systems can use these signals as a perception of the
>"weight" of the suitcase. So what we label "weight" is really
>"effort."

>Your proposal is another argument for the modified model we've
>been kicking around for a year or so, or maybe another related
>version of it. Can you draw if diagram of just how the various
>signals would be treated?

Would the perception of "weight" separate out the disturbance (e.g. someone putting their foot onto the suitcase as you start to lift)?

I'm hoping that a sensible diagram comes out of the "Information" paper, but in the meantime I'll see if I can draw up an intuitive one that might work. I doubt it would be very different from one you might draw, given our previous discussions. But we do need a straw diagram.

Martin

Date: Thu Feb 18, 1993 4:42 pm PST
Subject: Re: meaning, control, prediction

[Martin Taylor 930218 19:20] Rick Marken 930218.0900

>I thought that someone had claimed that control is only possible
>when the effect of the output on the input is perfectly predictable.

Well, that someone wasn't me, and I don't remember seeing that claim. My claim is twofold: (1) if the effect of the output on the input is perfectly predictable, then control (in the sense of a closed feedback loop, not Hans Blom's sense) is unnecessary; (2) if the effect of the output on the input is totally unpredictable, control is impossible. Closed loop control is both possible and necessary when the effect of the output on the input is partly predictable. In the tuning (focussing) case, that is exactly the situation. You don't know exactly how much to turn the knob, and there is initially a one-bit uncertainty as to which way to turn the knob (that is, indeed, all the uncertainty there could be about that question). But having resolve the one-bit uncertainty, one expects the same directional linkage to be maintained. Having moved the knob a little and seen very little effect, one is happy that moving it a little more is unlikely to make a big difference, and moreover is unlikely to make an opposite difference. That's partial predictability. Perfect for a control system (PCT sense).

I tried to avoid getting into the e-coli discussion, because it mixes up two things, the predictability of the environment and the predictability of the result of one move. The e-coli solution to control would not work if the gradient kept shifting around so that a move that made matters worse would, if continued make matters better in a couple of microseconds. E-coli depends on the predictability of the world, even though it does not control which direction it is going to move on the next tumble. Changing the rate would not work if the perception kept changing randomly. Likewise, an e-coli approach to the focussing problem would not work if the slide kept snapping back and forth at random, fast with respect to the bandwidth of your output actions.

>Just another wonderful benefit of the fact that control is the
>control of PERCEPTION, not OUTPUT.

Yes, that's my whole point.

Are we converging?   Martin

Date: Fri Feb 19, 1993 12:16 am PST
Subject: Re: meaning, control, prediction

[From Oded Maler (930219.0900-ET)]   Rick Marken (930218.0900)

You say, correctly (that is, what your words evoke inside me is coherent with my exoerience) :

```
* Law, religion, and other human foibles are based on
* the idea that words have meaning "in themselves"; why else
* would lawyers constantly argue about what sentences (laws)
* "really" mean; why would religious idiots try to explain
* what sentences (god's word) "really" mean. If people could
* be disabused of this transparently idiotic idea (that words
                                     ^^^^^^^^^^^^^^^^^^^^^^^^
* have meaning) and be educated to understand that words only
^^^^^^^^^^^^^^
* "mean" the experiences that they evoke in each individual --
* experiences that are likely to be considerably different across
* individuals -- maybe people could start approaching the problem
* of communication from a whole new perspective; one aimed at
* convergence on common experience rather than common wording.
```

and later you demonstrate how hard it is to educate people to get rid of this idea:

```
* Just another wonderful benefit of the fact that control is the
                                        ^^^^      ^^^^^^^ !!
* control of PERCEPTION, not OUTPUT.
```

wlsefno sfdqer   --Oded

Date: Fri Feb 19, 1993 8:00 am PST
Subject: Misc subjects

[From Bill Powers (930219.0730)]

Almost caught up.

Martin Taylor (930216 16:20) -- Hans Blom (930216) --

RE: nonlinearities in control systems.

Can't remember if I posted on this or not, except to Greg Williams last summer, but anyway --

Hans Blom said "...If the plant to be controlled has a non-linear but monotonic input-output transfer function, it can be controlled by a linear controller. It is just as if the system's loop gain changes depending upon the point of operation."

There's an interesting nonlinearity in muscle response, in that the force-length spring relation follows a pretty good quadratic curve over most of the operating range. When two such muscles are used in opposition by a set of balanced ("push-pull") control systems, an interesting result occurs (not yet put in the Arm model, but planned).

Consider a balanced pair of muscles in equilibrium. The differential tension in each muscle under passive stretch follows the law

f = k*(lo - l)^2 where lo is the resting length

If a pair of muscles is opposed, the net force developed for a displacement l from the equilibrium position lo is

F = f1 - f2 = k*[(lo - l)^2 - (lo + l)^2] or

F = -k*(2*l*lo)

The resting length lo is set by an equal muscle contraction on both sides of the pair. The differential displacement l is set by the difference in muscle contractions on the two sides.

Now suppose that there are two control systems, one controlling the sum of the forces and the other controlling the difference of forces. The sum controller sends the same output signal to both sides; the difference signal sends a positive signal to one side and a negative signal to the other side. Thus the sum and difference of forces can be controlled independently, with respect to independent reference levels. The combination is equivalent to two actuating signals, one larger than the other but both positive.

The sum-of-forces system controls muscle tone; the difference-of- forces system controls the net force tending to accelerate the limb about the joint.

More to the point, the sum-of-forces system varies lo while the difference-of-forces system varies l. If the sum is held constant, the spring constant with respect to difference signals is now linear. And as the sum of forces is varied, the net spring constant varies in the same way. So the sum-of-forces system controls the stiffness of the LINEAR spring for the difference- of-forces system. This will alter the frequency-gain characteristics of the control system that moves the limb, which in turn can help maintain stability over a range of load masses.

This seems to contradict Hans' statement

>My point was, however, that as soon as ANY type of non-linearity
>exists, only non-linear control schemes will lead to peak performance.

A contradiction ... unless the control of the sum of forces is counted as a control scheme (linear) having a nonlinear effect (variation of the spring constant for the difference-control system). Both of the actual control systems are linear, however.

Hans Blom:

```
> ...You can do three things with a model:
>
>      input     -------------   output
>      ----------|   model   |---------
>                -------------
>
```

>1. given input and model, we can calculate the (most likely)
>outputs; this is called prediction;
>2. given model and output, we can calculate the (best) inputs;
>this is called control;
>3. given input and output, we can calculate the (best) model;
>this is called system identification.

Martin assumed that you were talking about the environment here, but I'll assume that you're talking about an active system. While this arrangement is verbally neat, (2) doesn't make much sense to me because it omits any closed loop.

If the "model" is a closed-loop control system, it is really of this form:

```
      input     -------------   output
      --------->|+  model   |->-------
              ->|-          |  |
             |   -----------   |
              -------<---------
                 feedback path
```

In that case, the "input" is not a sensory input, but a reference input. The sensory input to the model is located where the feedback path branches off the output path.

If you map this diagram onto a diagram of, say, a spinal reflex, you find that inside the model box there is a comparator in the form of a spinal motor neuron excited by the "input" signal and (for the tendon reflex) inhibited by the feedback signal. The error signal enters a muscle, the actuator, which produces the "output" force. The output force is sensed by Golgi tendon receptors, which produce the feedback signal.

The "input" signal in the real system does not come from the environment via a sensor. It is instead a signal that descends in the spinal cord from centers higher in the brain. It is not, in fact, an input to the organism, but an output of higher systems. The only sensory input is the tendon receptor that senses the output force.

In a hierarchical control model, no reference signal comes from sensory receptors; all are either the outputs of higher control systems or are genetically fixed. Organisms, therefore, do not have "inputs" in the sense shown in your diagram, according to the HPCT model. All control systems actually traced out in the nervous system are of this kind.

I would like to point out another difficulty with point (2) in your list above. There is in fact a closed loop even when there is no feedback inside the model. When you speak of finding the "best" input to produce the "given" output, there is implied a reference level for the output -- the definition of "given". In order to find the best input, you (the designer) must sense the actual output, compare it with the desired ouput, and use the resulting error information as a basis for adjusting the input or the parameters of the model. So the designer becomes part of the closed loop: the sensors and comparator are in the designer.

In any actual non-feedback "control" system, the output will remain the desired function of the input only as long as all the components (including actuators) retain their calibration precisely. All real systems designed this way work perfectly only for a limited time. Then, if they are to continue working properly, the engineer standing by with his tools must twiddle the adjustments, while watching the actual performance and comparing it with the desired performance. So the continued functioning of an open-loop system depends on presence of a closed-loop system.

This is why I say that all real control is closed-loop.

Wolfgang Zocher (930218) --

SIMCON4.3 received. Will check it out today.

Gary Cziko (930217 I think) --

RE: Focus control.

I can't find the reference, but it has been found in experiments with the focus control system of the eye that when an image jumps out of focus, the initial change in eye focus is wrong 50% of the time.

So the control system involved, as Rick says, can't be simple. It must detect the relationship between the direction of effort and the direction of change of the absolute error signal, and if the relationship is wrong, flip the sign of the conversion from error to output. At least two ECS are required, one of them sensing a relationship and acting by reversing the feedback sign of the other.

Rick Marken (930218.0900) --

>>Of course real control engineers know a lot more than I do
>>about the design of complex control systems, and some day they
>>will take PCT much farther than I possibly could.

>No offense, Bill, but this sounds somewhat disingenuous.

No coyness intended. Real control engineers have mathematical skills that I do not have. I am often limited in my modeling by not knowing how to do the analysis of a system. Finding ways to do it costs me many hours, days, or weeks of sweat and I don't always succeed or find the best or simplest way. My main advantage over "real" control engineers is that I have developed a strong intuition about control processes that is simply not taught in engineering courses. Engineers tend to stick to conventional ways of seeing control systems, and thus miss concepts like control of input. The first real control engineer who grasps my intuitions will carry us forward by a long distance.

From Oded Maler (930218.1900-ET) --

>If you stay with circles, it might be interesting to have a
>"motor program" (your version of the concept, that is with
>continuous perceptual feed-back, no access to "objective"
>variables) that draws circles robustly ...

I guess this could be done with circles or other regular figures -- sensing the radius of the circle shouldn't be difficult. I'd be very happy if someone else worked this out, though.

Best to all, Bill P.

Date: Fri Feb 19, 1993 9:30 am PST
Subject: WTP LEVELS - RKC

[From Bob Clark (930219 11:45 am EST] Bill Powers (930216)

I was preparing for a final edit of "Higher Levels: II" when I received your post (above). I am pleased, but not surprised, to find our primary views of "the world" have remained identical over the years:

>This world, to the best of my knowledge, originates in signals emitted
>into the nervous system by sensory receptors.

And:

>This means that the world we experience must consist of sensory signals
>and other signals derived from them. The "other signals derived from
>them" include the totality of what we can experience, from the taste of
>chocolate to Fermat's Last Theorem, as well as our experienced "interest"
>in that Theorem, if any, and any "thoughts" we may have about it. Nothing
>is exempt.

Also:

>When I say "it's all perception" this is what I mean. We live inside a

>nervous system and all we know is what goes on inside that nervous system.

Given this viewpoint, with which I completely agree, there are several pertinent problems.

You report that your "pre" idea was, from paragraph 3 of your post:

>attempting to characterize human beings by identifying levels of control
>with various aspects of human functioning.

That does not quite fit my recollection, but we probably need not resolve the matter at this time.

My present views have developed irregularly over the years. They have been modified since you and I were in contact in 1987, and further developed since I met Greg Williams in 1988. Some of the ideas I have been presenting recently are still being revised. I certainly expect further changes as discussions proceed -- just as I think you also expect.

In your current post, you have re-stated your current view:

>... the problem in understanding human nature is not so much to understand
>human beings as to understand the world that human beings experience.

>This changes the problem. Now the problem is to classify all of
>experience, not just experiences of other people. We may perceive another
>person driving a screw into a piece of wood as showing a "skill" type of
>control, but this leaves unexplained the screwdriver, the screw, the piece
>of wood, and the relations among them.

See "++" below.

You use two familiar, frequently used, words: "understanding" and "explanation." Exactly what does each "really" mean? I find my dictionary of little help here -- let me try to define them:

"Explanation" seems to consist, at a minimum, of being classified, that is, placed in a category. That category may or may not pre-exist, but to be useful, probably should contain more than one element.

Is a dog "explained" by having its breed specified? Or naming its species? Or its genealogy? How about its physiology, or neural systems? Of course not. Neither is "control of a perceptual variable" "explained" by pointing out that its actions resemble those of a negative feedback system.

Instead of "explaining" some thing, activity, system, or whatnot, I prefer "description" of parts, their connections with each other and with other items. "Interactions" among the parts and with other items describes its "behavior." I am pretty sure that this is what you mean by "explanation" and "explaining."

"Understanding" is the goal of every teacher for his students. Me too. However it seems to me that there are two aspects to this concept: internal and external.

The "internal" aspect is displayed by simply asking, "Do you understand this matter?" If "Yes," is the reply, this signifies that there is no perceived recognition of inconsistency within his internal array of information (perhaps after modification to include the new material).

The "external" aspect is more complicated, being displayed by asking the other party to "solve" a problem that requires "proper" use of the material to obtain "the" solution. If the result is "acceptable," it indicates (does not "prove") that the comparable parts of each party's systems are in agreement. Desirable, of course, because further discussion is facilitated, possibly leading to revision (perhaps by both participants).

It is interesting that I have had the experience of saying "yes" to the question, but finding the external test reveals some degree of "misunderstanding." Indeed, I think that most people have had this experience in one form or another.

"Consistency" is demonstrated by this procedure, but not necessarily consistency with other parts of either party's systems.

++
"person driving a screw into a piece of wood as showing a "skill" type of control ... screwdriver, the screw, the piece of wood ... etc"

In my view "skill" is not a "type of control," rather it is a combination of perceptual variables that includes perception of objects (screwdriver, etc), the one using the tool, the location of the several objects, the sequence of events and interactions required in order to "drive a screw into a piece of wood." This "combination of perceptual variables" includes several less complex skills, such as reaching for the screw, placing it in the required position, etc. This entire combination may be referred to as "driving a screw etc," which is one among many muscle skills that may be used to accomplish higher order purposes.

Thus "Skill" is a category of perceptual variables, selected for purposes related to interactions with other people and distinguished from lower order variables by combining them (sequences of muscle tensions combined with temporal variables) to form the specific skill selected.

Perhaps that is not an "explanation," but I think it is "understandable" and I hope that it communicates something of my view of Fifth Order. ++

Time out to attend a meeting of the Dayton Chapter of the American Helicopter Society. 50 miles each way to Dayton. A rather interesting presentation of some of the current developments with the Tilt-Rotor helicopter. I will turn back to "Higher Levels: II" after I post this, and review my Inbox.

Regards, Bob Clark

Date: Fri Feb 19, 1993 9:55 am PST
Subject: belated insight;consciousness; measuring higher cps

[From Dick Robertson] (930219)
Subjects: (1) insight about controlling output; consciousness;
(2) investigating higher levels

After I hung up yesterday I felt really stupid about my question about how output is "controlled." I realized that the output signal rises as the error rises - the whole question of "computing output" was a relapse on my part into the old "people control their outputs" thinking. But why did I have to ask my stupid question before the answer came to me? Why didn't I work it out in my head when it first came up?

This reminded my of something Daniel Dennett talks about in Consciousness Explained. (Did someone mention him a few months back, or am I imagining that?) He proposes that our thinking consists of "multiple drafts" in which different control systems (my terminology, not his) in the brain compete for consciousness...and therefore we often don't know what we think until we hear ourselves say something. I thought of it in terms of different versions of a program competing to satisfy some error in a principle-level system. It's a terrible shame Dennett doesn't know control theory; his book would be much more exciting, if he did, I think, but even so there is a lot of interesting speculation in it.

For instance, some of the research he reviews seemed consistent with one of my favorite speculations about consciousness - that it has something to do with shifting perceptions to areas in which errors are being corrected. For example, my attention shifts to bring more information (i. e. to reset lower order RSs for correcting higher level error signals) - e.g. to body- orientation-sense when I slip while walking, to the road when I'm daydreaming in the car and there is a traffic glitch ahead, or to the placement of my fingers on the keyboard when wrong letters start to appear on the screen, etc.. In that last instance it directs macular vision to the hands when peripheral vision isn't enough to keep the hands from straying to the wrong position. We often informally call consciousness a "spotlight" as simply a metaphor. But it strikes me that at least with vision it seems litterally correct - my visual awareness consists of what my macular vision is focused on as compared with all the additional information that is coming in around the edges. I can inwardly shift my attention to peripheral information deliberately, but the difference between that kind of consciousness and macular consciousness is the difference between recalling someone's face in memory and seeing him directly.

Are there auditory and tactile counterparts of macular vision? Does anyone know?

....

I want to bring those who heard my report on the students' grade control study at last year's CSGconf up to date. I have now run the study three times and ready to give up on the current approach. I am convinced that some students do control their grades literally, others do not. The latter seem to approach their course grades as like a lottery; they make a standard effort and take whatever they get. But even for those who do control their grade__as evidenced by remarks that they had to decide to study harder than they had planned to, or had to give up activities they had intended in order to study__the method by which I tried to measure it didn't pan out. For the last sample I got 4 subjective ratings (5-pt scale): (1) how hard they studied - rated before the quiz; (2) after the quiz - how hard it was; (3) how well they understood the material; (4) how many they thought they would get right. None of these measures are well related with either the score on a given exam or that on the just-prior exam for any of the students I have looked at so far, nor does a composite of the first 3 judgments.

But, I don't feel like giving up on the problem. I can't believe that measuring higher order control is that hard. Since my self-image-control measure keeps replicating reliably I feel that there must be some real- time, immediate grade control efforts which are comparable to tracking performances on the one hand or self-image control on the other. In both of those cases the key to observing on-going control efforts seems to be the immediacy of the observation. I can't figure out what would permit a similar immediate observation of whatever grade-control is going on. The results I got reminded me of stuff I learned years ago about the problems of psychological test construction, namely that instead of simply marking items (on the MMPI, e.g.) "honestly" subjects would mark them in terms of what they thought would make a good picture of themselves, or what they thought were "socially desirable" traits to have, etc. I can imagine that students try to control various images like that in doing the ratings__even were they able to compare the amount of study they did on one occasion with that of another.

How would one parse what's really going on here to set up a model the way that the tracking experiments are simplified analogies of what one does in steering a car, e.g. Would it be possible to take a task like the old "Powers game" and interrupt it periodically with other tasks, so that we could see a subject shifting control from one perceptual variable to another. Have you, Bill P, Rick M or Tom B already got a task in which that happens?

Best, Dick Robertson urrobert@UXA.ECN.BGU.EDU

Date: Fri Feb 19, 1993 10:08 am PST
Subject: Re: meaning, control, prediction

[From Rick Marken (930219.0800)] Dick Robertson (930218?) --

A hard carriage return or two dropped out of your post with the
question for me -- so I could not understand it. But it seemed
like it was going to be very interesting. Could you try posting
your question again? Thanks

Martin Taylor (930218 19:20)

>E-coli depends on the predictability of the world, even though it does not
>control which direction it is going to move on the next tumble. Changing the
>rate would not work if the perception kept changing randomly.

I have to disagree with this. One virtue of a control system organization is that it is able to produce predictable results for itself (its perceptions) in an unpredictable environment (and, by environment, I include everything outside of the nervous system). The e. coli system demonstrates this fact in spades; in fact, the perceptions experienced by e. coli as a result of it's actions could be completely random (equally probabale) yet the control organization would be able to STOP acting once it was experiencing the intended perception -- and, thus, keep experiencing that perception. So e. coli CAN control the results of it's actions even if those results are perfectly unpredictable.

>Are we converging?

Certainly. I'm just being tenacious about this because I think you are trying to say that statistical properties of the environment place constraints on the nature (or existence) of control systems; or something like that. I think this is just another case of trying to preserve some beloved concepts from the past. There is nothing wrong with beloved old concepts, per se, but its best if they have some demonstrable value. I may be misunderstanding your point about "predictability" (because it is evoking images in me that are not the same as those it evokes in you). I think the only way to solve this is through demonstration and modelling. I can show you how an e. coli system (and a person in the same situation) can produce a particular intended result even though every possible percetual result of action is equally probable. Could you invent a demonstration that would help me understand your point about predictability?

While we're on the subject of meaning, I said:

>If people could be disabused of this transparently idiotic idea (that words
>have meaning) and be educated to understand that words only
>"mean" the experiences that they evoke in each individual --

and Oded Maler (930219.0900-ET) says:

>and later you demonstrate how hard it is to educate people to get
>rid of this idea:

* Just another wonderful benefit of the fact that control is the
                                        ^^^^   ^^^^^^^ !!
* control of PERCEPTION, not OUTPUT.

I don't get it? Are you saying that this statement, in particular, reflects a belief on my part that words themselves have meanings? Is it the word "fact" that causes a problem? Do you imagine that the "evoked meaning" view of words rules out the idea that there are "facts"? I didn't mean to imply "ultimate reality" here. The fact that control is the control of perception is a fact in the sense that it is an experience that you can reliably produce for yourself in various ways -- through demonstrations (like those described in "Mind Readings" and Bill Powers' Demo program) and mathematical analysis. It is a much a fact as the "fact" that electric current flows from a point of high to a point of low voltage, the rate being proportional to the voltage difference and inversely proportional to the resistance of the medium.

Best   Rick

Date: Fri Feb 19, 1993 11:48 am PST
Subject: Re: meaning, control, prediction

[Martin Taylor 930219 14:10]   Rick Marken 930219.0800

> One virtue of a control system organization
>is that it is able to produce predictable results for itself (its
>perceptions) in an unpredictable environment (and, by environment, I
>include everything outside of the nervous system). The e. coli system
>demonstrates this fact in spades; in fact, the perceptions experienced

>by e. coli as a result of it's actions could be completely random (equally
>probabale) yet the control organization would be able to STOP acting
>once it was experiencing the intended perception -- and, thus, keep
>experiencing that perception. So e. coli CAN control the results
>of it's actions even if those results are perfectly unpredictable.

But there's precisely the point. In an unpredictable environment, if it stopped, it would not keep experinecing that perception. It is the partial predictability of the environment that permits it to succeed.

E-coli may not be able to imagine what the result is going to be of its next move, but it does work on the principle that if the move makes the perception further from the reference, then another move is in order. In an unpredictable environment, this wouldn't work.

>I'm just being tenacious about this because I think
>you are trying to say that statistical properties of the environment
>place constraints on the nature (or existence) of control systems;
>or something like that.

There's no "something like that" about it. This is exactly what I am trying to say.

> I may be misunderstanding your point about "predictability" (because it
>is evoking images in me that are not the same as those it evokes in you)

I'm pretty sure that this is the case, which is why I have been rather long-winded and pedantic about the issue.

> I can show you how an e. coli system (and a person in the same situation)
>can produce a particular intended result even though every possible percetual
>result of action is equally probable.

I know you can, and classical (Powers) reorganization works this way, too. And both work very well. What you forget is the control of the rate of movement is the critical way that e-coli and reorganization work, and in an unpredictable environment, that could not work.

Do your e-coli demonstration in an environment where the gradients change randomly at every time step, and the perception at any point is represented by a random value at every time step. Then see if it can control.

>I'm just being tenacious about this...

Me, too. Unless the important truths lie in a different dimension, one or other of us will change our understanding eventually. It's best that we keep problems like this overt, rather than allowing the underlying misunderstanding to spoil communication on other topics.

Martin

Date: Fri Feb 19, 1993 1:28 pm PST
Subject: social convention

[Bruce Nevin (930210, 930218-19)]

My participation here will be limited, for reasons that I have suggested -- recent RIF, promise of more, new requirement that I work 9-hour days, high visibility of any lapses of attention, etc. A number of threads will lapse, I am afraid. Martin, with you in particular I am aware of numerous unclosed loops, and I am sorry that they will remain so indefinitely, but that's the way of it.

I have been out sick for a week with a vicious case of the flu. I started writing this at home last week. Then I didn't feel like doing anything for a great while. I took it up again last night (2/18 evening), and finished this morning. I may yet go in to work, as my boss is worried about HR's deadline for four performance reviews, one of them mine (I have to add my comments and sign off), and we have to pack our offices for another move on Tuesday. (The floor we were on will be sublet.) I will see if I can break this up into several instalments.

The thread that most concerns me of course is the question Bill and I were beginning to address:

(Bill Powers (930204.1500)) --

>What I'm most interested in is what a social convention (for anything)
>is that we can perceive it, and what it is we must perceive in order to
>know that we're experiencing one, or conforming to one.

Most of the time that we are conforming to social conventions we don't know that we're experiencing them or that we're conforming to them, and we don't need to. The social conventions that we become aware of are almost always shibboleths.

The term is from Judges 11, a tale worthy of Grimm's. The part relevant to us here is where the Ephraimites were ticked off at Jephthah and his Gileadites for various reasons. They fought. The Gileadites held the fords of the Jordan. Anyone who wanted to cross, they asked if he was an Ephraimite. If he said "no" they commanded him to "say `Shibboleth'," which means "running stream". The Ephraimites pronounced it `Sibboleth', so those posted at the fords killed them on the spot.

The term is used today in reference to the self-betrayals of behavior by which social-class distinctions are guarded. Much of instruction in "manners" concerns avoiding ways of doing things in the manner of one social class in favor of doing them in the manner of another.

Most of the tokens of membership in a social group are not accessible to self-conscious manipulation. The reasons for this, and the fact of it, we appear to hold in common with all mammals (according to Bateson), certainly higher mammals. To see this, we need only recognize that the "unconscious shibboleths" by which group members identify one another and identify with one another are only a special case of nonverbal communication or "body language." The arguments I advanced here quite some time ago about nonverbal communication apply a fortiori to the special case of group membership -- that it must be unmanipulated to be accepted as authentic, that we are very quick to detect inconsistencies that betray the taint of inauthenticity, that we profoundly distrust politicians, salesmen, and other "two-faced" persons. (An aside: I don't think the proponents of NLP, etc. realize how profoundly they may be eroding social relations by putting sophisticated tools of persona- manipulation into the hands of legions of resolutely shallow people.) I will not repeat those arguments now.

A second matter that I would like to include by reference, if I may, is the notion that there is "free play" in the _manner_ in which one controls for one or another goal. This is really nothing more than an application of "Guthrie's lens" (described in BCP), the idea, familiar to CSG readers, that an indefinite variety of actions may serve the attainment of a goal. The application is this, that among these differences that don't make any difference for reaching the goal (pick up the cup in one's fist or with one's fingers on the handle, you get the drink either way), there may be some that make a difference for some goal at a higher level, associated with a manner or style. The higher level goal is demonstrating membership in a given social group. Or rejecting it, by flouting a stigmatized manner. Indifference is also possible, but not likely. We do not become human without becoming social beings, in the sense of "owning" a public face and identifying with it as our very selves; those who never become adept at membership, through deprivation (Genie) or other reasons, are clearly not fully human, in some important sense.

OK, now let's look at an example that I have introduced before, and then consider how we might model what is going on here. Natives of Martha's Vineyard Island speak a dialect that preserves the pronunciation of r after vowels, in words like barn and farm, whereas these words are pronounced r-less throughout most of New England. In addition, the first vowel of the diphthong in words like "about" and "house" is more centralized than in the speech of "summer people" from the mainland. In a generalized New England dialect, the first vowel of such diphthongs is more like the [a] of "father," on the Vineyard it is more like the vowel [A] heard in "up" and "love." (I am writing [A] here for the inverted v character that represents this sound in the International Phonetic Alphabet.) The tongue is higher, giving the au diphthong a kind of "clipped" [Au] quality that is also heard in various Canadian dialects. These characteristics were the norm in 17th-century English. The change came to Boston from certain dialects of London, and spread thence. The change seems to have been motivated by emulation of wealthy and fashionable people. The older pronunciations were preserved only in a few conservative pockets scattered around New England.

We'll consider only the diphthong. Towards the close of the 19th century, dialect surveys show that the older pronunciation had been fading, even on the Vineyard, in favor of the more open [a] sound that prevails elsewhere. But then this change began to reverse itself, just about the time the Vineyard Economy, no longer sustained by whaling (and secondarily fishing, brick-making, and sheep-raising), began its transition to a resort community dependent on "summer people". Surprisingly, the centralized [A] vowel quality came to be heard also in the other common diphthong, in words like "slide" and "wife", apparently by analogy.

It turns out (as reported in Bill Labov's thesis, and a paper published subsequently) that the process by which this happens appears to hinge on the choice of social identity by adolescents. Those who plan to leave the Island for better economic prospects on the mainland take on a New England or New Yawk dialect; those who elect to make their lives on the Vineyard begin speaking a Vineyard dialect. Speaking in one manner or the other is an unconscious declaration of affiliation and intentions.

The kids are not aware of changing their speech, certainly not of changing particular aspects of it. My oldest daughter, then 16, seemed unaware of taking on a Gloucester working-class accent, according to her chosen affiliations when she moved here from Seattle. Adamantly unaware, I might be tempted to say. Of course: as noted re mammalian communication, it is important that these differences not be accessible to conscious manipulation. Some of the intense self-consciousness in this "age of perpetual embarassment" is I think due to the intimately felt dilemma of having to do just this--manipulate body language, with all the inauthenticity and inappropriateness that this entails--in order to create a socially viable identity or persona that can be a vehicle for authentically representing one's intentions and expectations to others, and eliciting cooperation from them not as a child but as an adult.

I think this is one reason why, in these realignments of body language, adolescents famously tend to overshoot the mark. Thus, the adolescents who were staying assumed a Vineyard dialect even more conservative than that of the lobster fishermen of Menemsha and Chilmark, and those going assumed a corollary exaggeration of mainland dialectal traits.

In this are a few suggestions, if we are willing and able to take them up and examine them carefully, of how and when and why we learn social conventions, and of what they are.

>. . . this must involve noticing patterns in what other
>people do and also experiencing the results of doing things differently.

It must surely involving noticing at the time of learning, but (under the overriding stipulation that things of this kind should not be noticeable, that is, cannot subject to conscious manipulation if they are to be authentic) thereafter the comparison is not made to the behavioral outputs of others, but to the internalized reference perceptions that one has learned. This makes "overshooting the mark," so characteristic of adolescence, more likely. (Small children, I take it, do not feel this dilemma anywhere nearly as acutely. Their "overshooting" is overgeneralization of some newly learned pattern in language, not unconscious caricature.)

> Clearly we don't imitate everything other people do;
>otherwise we'd all be doing exactly the same things. How do we
>choose which patterns to imitate and which to do our own way?

These questions have been answered above, I think?

>What is it that we do when we see someone else deviating from a
>pattern we have accepted as the socially right one?

Depends. We might perceive them as "misbehaving" if it's a shibboleth. Otherwise, we would focus on goals that matter to us, such as successful cooperation with them, and overlook the differences that don't make any difference for those goals. If the discrepancies are recurrent, numerous enough, and systematic, we may perceive them as "foreign". We may begin to take on some of the discrepant mannerisms ourselves (I'm a regular chameleon, I'm told), for the sake of strengthening bonds of commonality between us, or if we perceive them as members of a social class to which we would like to claim or pretend membership, or for various other reasons.

In a justly famous 1923 article, "The psychological reality of phonemes," Edward Sapir gives a number of striking examples of how the patterning in language is the basis for overriding perceptual input of a given sound-token in favor of imagined perception of an exemplar of the sound type (category) to which the token belongs. I will describe the last of these acoustic illusions briefly; it is the most accessible, since it depends upon facts about English.

Every year that Sapir taught a course in "practical phonetics," he found that as soon as he had taught his (English-speaking) students to recognize the glottal stop ['], they frequently heard and transcribed it "after a word ending in an accented short vowel of clear timbre (e.g., a, E, e, i)," when in fact he had not produced a glottal stop there in his dictation.

(Here, I am writing E for Sapir's epsilon, representing the vowel of "bed", as opposed to the higher [e] of "attache." The glottal stop, recall, is the medial consonant in "uh-uh" ['A'A] meaning "no" and "oh- oh" ['A'o], with [A] the vowel of "up" as previously.)

Here is the pattern in the language that is the evident basis for the illusion. The following kinds of accented final syllables occur in English:

A. Words ending in a long vowel or diphthong: sea, flow, shoe, review, apply.

B. Words ending in a long vowel or diphthong plus one or more consonants: ball, cease, dream, alcove, amount.

C. Words ending in a short vowel plus one or more consonants: back, fill, come, remit, object.

There is a theoretically possible fourth class that does not occur in English:

D. Words ending in a short vowel: French fait, ami, attache; Russian xarasho.

English pronunciation of foreign words of type D is "drawled," making them words of type A.

Now, let's look at the transcription errors committed by Sapir's students when he would dictate nonsense syllables of type D, like [smE] and [pila] (accent on the last syllable). On the basis of the familiar sound pattern of English, they might "legitimize" the unfamiliar final accented short vowel by either of two illusions: They might imagine the perception of length, writing [smE:], [pila:] in their phonetic transcriptions (class A), though the vowel was indeed short. Or they might imagine the perception of some consonant X after the accented short final vowel (class C). The glottal stop, newly added to their repertoire of possible consonants, solves the problem for X, and they write [smE'], [pila']. "The error of hearing a glottal stop where there is none, in words of type D, is fundamentally a more sophisticated form of the same error as hearing a dictated final glottal stop as p or t or k, which occurs frequently in an earlier stage of the acquiring of a phonetic technique."

Sapir's example of English-speaking students of phonetics hearing a glottal stop where there was none hinges upon what is a possible syllable in a given language. Native speakers of English have certain expectations as to what is possible and what is not, exemplified by the little corner of the sound pattern of English that we have just lifted and turned about in the light. A native speaker of French would never hear a glottal stop after a short accented final vowel, given the different sound pattern of French. (French students of phonetics have their own difficulties with glottal stop, let it be said, for entirely different reasons.) Now this sort of pattern is, I think, not "a social convention" in the sense that placing forks on the left in a place setting is a social convention. It is an observable property of the body of words in use round about as a child learns the language. The child develops a sense of pattern as a basis for setting expectations as to what might be being said.

There is nothing that constrained us, in our nature as individual control systems, to develop the particular pattern of English as we grew out of infancy, as opposed to that of French, or Hawai'ian, or (God help us) Bella Coola with its clusters of six or more consonants. To be sure, there are constraints complicit with control of the acoustic, kinesthetic, tactile perceptions involved in speech, but within those universal human limits further constraint is determined by what the child finds already present in the speech community, as historically contingent social fact.

To model this, set high gain on control for the perception of participation or cooperation. One form of cooperation with obvious utility is when an adult serves the child's needs by means that are beyond the child's present capacities. Another is when an adult guides the child through some scaled-down "play" version of adult skills and practices. I wonder if there has been much thought put to the child's perception of participation? It seems obvious to me that very much of children's clamor for attention amounts to unskilful bids for participation and inclusion.

Alongside this, I suspect there is strong motivation (intrinsic?) for the development of higher levels of control. The acquisition of language must be synchronous with development of control above the event and relationship levels. Learning to participate with adults using language has then the double motivation of participation, pure and simple, and learning "the world" as it is socially organized by our talk about it.

And here is another reason why the inherited, learned patterning of language and its organization of perception is not readily recognized as conventional: it is seen as the way the world is, and foreigners are simply mistaken.

Before I close this long review. You will recall, Bill, the discussion that left you with the giddy sensation of the floor and ceiling being taken away? I want to reach beyond the comfort zone just briefly now, with the observation that unconscious social convention for the sake of effective participation as a member of a group or community fits extremely well into the profile that we drew of how control systems of a given *order* (e.g. cells) must implement (and yet not experience per se) control systems of the next higher order. Following this conjecture, if we wish to understand how reorganization takes place among cells, perhaps we should look around. Individually, we may even be able to participate consciously in the process, and to develop some skill in it. That, in my view, is the proper perspective to take on clinical applications of PCT.

- end -

Bruce Nevin bn@bbn.com

Date: Fri Feb 19, 1993 1:43 pm PST
Subject: HIGHER LEVELS: II - RKC

[From Bob Clark (930219 16:00 EST)] Bill Powers (930201.1900)

>To speak of "personality" and "character" is to take an external view of
>someone else's organization. That is, you seem to be looking for levels
>that will apply to "psychological" aspects of a person, to explain the how
>and why of that person's behavior.

I included a few comments on this item at the end of my post of "Leader/Follower" (930206.2212 GMT), and also at the end of "Higher Levels: I" (930215.1209 EST).

Contrary to your suggestion, my analysis is based on your very important observation that: BEHAVIOR IS THE CONTROL OF PERCEPTION, with perceivable variables being the basis of the structure.

Additions and revisions are certainly needed. A great deal remains to be done!

Since our views of the lower levels are rather similar (with the possible exception of my 4th Order, Temporal Variables), we move to higher levels.

Here I seek controllable, perceivable variables that are formed by combining lower order variables. It occurred to me that muscle skills can be regarded as sequences combined with temporal variables. There are many such perceivable combinations. Some are relatively "simple," like walking, pressing fingers on buttons, pulling rubber bands, etc. And some are very complex skills like vocalizing, running, throwing, dancing, acrobatics. Thus Muscle Skills, a group of perceivable, controllable combinations, can be assigned to Fifth Order, "Skills."

Such muscle skills are readily perceived not only in others, but also in oneself. Many are learned, some probably have genetic origins. In the process of learning how and when to use them, variations of many sorts are explored. Such experiments and their results are recorded (as "memories") as they occur. Thus they remain generally available for later use.

What comes next? What would be the nature of Sixth Order Activities composed of controllable, perceivable variables based on combinations of lower level variables, especially Skills of Fifth Order? As I was seeking to distinguish Fifth Order from Fourth Order, there was a tendency to consider interactions between/among individuals. Thus, with Fifth Order assigned to "Skills," Sixth Order could include all activities using combinations of Skills for purposes requiring control of interpersonal interactions. Examples include games, competition, cooperation, government, clubs, businesses, entertainment. In addition, language, mathematics, philosophy, systems, principles, programs are included here. Here we find all theories, whether of the natural world, the world of imagination, the world of behavior, etc, including Perceptual Control Theory.

People generally have some sort of structured view of the nature of their surroundings and how to achieve their objectives. Their methods may be based on gross misunderstanding, superstition, or what, but they are sufficient for most people most of the time.

Communication, complex combinations of many muscle skills, taking many forms, is used throughout interpersonal interactions for many purposes. Should this be considered another level?

In examining that possibility, it occurred to me to pay attention to everyday conversations among my friends and associates. Much conversation pertains to Zero-Order systems -- health, sensations of temperature, physiological events. There was discussion of combinations of sensations perceived as "objects." In turn, sequences forming postures, movements, etc were of interest. These various combinations were used for ordinary, customary purposes of communication.

As "Topics of Communication," these may be called "Modes" of Sixth Order, corresponding to Orders of control, without themselves being control systems. Topics relating to Skills would be Fifth Mode of Sixth Order. Those relating to Communication and other interpersonal variables would be Sixth Mode of Sixth Order. The Modes do not function as Control Systems, but assist in analyzing the structure and performance of the Systems.

Continuing these observations, one finds comments about personalities and characters of individuals. What does this mean in terms of perceivable variables? The dictionary answers these questions rather well.

"Personality: 1. The visible aspect of one's character, as it impresses others: "He has a pleasing personality.""

This looks as though it could belong to Sixth Mode of Sixth Order, but it seems to me to go a bit further. Thus we have people who are actors, behaving to portray varying personalities, emotions, etc. They appear to be controlling their behavior to produce certain interpretations by those around them. Being "pleasing," "friendly," "courteous," "hateful," whatever, can be controlled, even if contrary to the performer's own internal feelings. Thus "personality variables" can be regarded as controllable, perceivable variables in the performer's own repertoire. Interestingly, because combinations of skills are needed to display these variables, the time scale needed to perceive these variables is moderately long vs the time needed for demonstrating lower Modes.

"Character: 3. moral or ethical quality, 4. qualities of honesty, courage, or the like; integrity."

Other definitions seem too inclusive or specialized. I think this does pretty well. Here there is another increase in the time scale. While personality can sometimes be demonstrated in a matter of minutes, character requires observation of several incidents distributed over a much longer period.

These topics, "Personality" and "Character" are sufficiently different from each other and the other Modes of Sixth Order that they could be treated as Seventh and Eighth Modes of Sixth Order. Their importance in forming "images" of other people also suggests assigning them to Seventh and Eighth Modes of Sixth Order.

This assignment would imply the existence of Seventh and Eighth Order Control Systems, based on corresponding perceptual variables.

This discussion suggests that something like "Self-Image" could be considered Ninth Mode of Sixth Order, with corresponding Ninth Order Control System. This treats Personality and Character as important components of Self-Image, in addition to all other perceptions of whatever composes one's "self."

Where and how the DME, "Decision Making Entity," would relate to this structure is postponed for the present.

In my post of December 5, 1992, (921205), I suggested "Self-Image" as Seventh Order, with a corresponding Mode of Sixth Order. At that time I was not yet satisfied with the distinctions among Personality, Character, and Self-Image. I am still not very confident that these are appropriate, but they may be useful for discussion.

Conceived, I think, as a truly general theory of behavior, Perceptual Control Theory should apply not only to observations of the behavior of other people, but also to ourselves, both individually and in the process of constructing a Theory of Behavior.

As suggested in the post quoted at the beginning of these remarks, "Personality" and "character" certainly can be used for describing other people. The remark about

>"looking for levels that will apply to "psychological" aspects of a
>person, to explain the how and why of that person's behavior"

seems to imply that my suggestions do not apply to oneself. The only basis for that implication seems to be the observation, with which I agree, that: "it is impossible to perceive your own behavior while it is occurring." However this can be accomplished by examining your own recordings of your own behavior. While this is not immediate feedback, it can be fast enough for many purposes. For these higher order systems, the time scale can be rather long -- days, years?

On examining my memories of my own behavior, I find I can generally perceive even these high order variables in my own remembered behavior. Perhaps more important, I find that, if I care to, I can generally change my behavior. This may take more time than I like, but my perceived and changed behavior has become more nearly what I sought.

Further revisions are certainly needed. Perhaps most important, PCT should be applied to problems of general interest.

Regards, Bob Clark

Date: Fri Feb 19, 1993 4:07 pm PST
Subject: Vonnegut quote

[From: Bruce Nevin (Fri 930219 18:40:34)]

Someone (Oded?) wanted the correct version of this quote. I, too, was struck by it and wrote it down years ago. Cleaning my office for the move, I came across a file card with this on it:

> Beware of the man who works hard to learn something, learns it, and finds himself no wiser than before. He is full of murderous resentment of people who are ignorant without having come by their ignorance the hard way.
>
> --Bokonon
> (In Kurt Vonnegut novel, _Cat's Cradle_ I think)

Date: Fri Feb 19, 1993 4:16 pm PST
Subject: Re: prediction, clinician

[From Rick Marken (930219.1500)] Martin Taylor (930219 14:10)

>Do your e-coli demonstration in an environment where the gradients change
>randomly at every time step, and the perception at any point is represented
>by a random value at every time step. Then see if it can control.

Ok. I think what you might be saying is that there may be environments in which control is not possible, period. This is what you have described above for e coli: at every resolvable instant in time the input is completely random -- NO MATTER WHAT OUTPUTS ARE GENERATED. I don't know if there is any real environment that has these properties but I agree that, conceptually, nothing could live (control ANYTHING) in such an environment.

My original point was only that control does not depend on even partial predicability of the effect of output on input. The situation you describe seems less like a predictability problem than a bandwidth of disturbance problem -- the disturbance (or whatever you want to call the cause of random variation in the input) is occurring at a rate that exceeds the mimimum time resolution of the system.

Rather than "a partially predictable environment" I might be able to go along with the constraint that control systems cannot exist in an environment where the only disturbance is completely random and of infinite frequency.

Bruce Nevin (930210, 930218-19) D --

Very interesting and challenging series. But your punchy finish really caught my interest:

> I want to reach beyond the comfort zone just briefly now,
>with the observation that unconscious social convention for the sake of
>effective participation as a member of a group or community fits
>extremely well into the profile that we drew of how control systems of a
>given *order* (e.g. cells) must implement (and yet not experience per
se) control systems of the next higher order. Following this
>conjecture, if we wish to understand how reorganization takes place
>among cells, perhaps we should look around. Individually, we may even
>be able to participate consciously in the process, and to develop some
>skill in it. That, in my view, is the proper perspective to take on

>clinical applications of PCT.

I believe "going up a level" is the essence of the clinical approach of PCT; but I think you are talking about some rather different levels than what I understand as the PCT levels. When I "go up a level" the experience is quite "mundane" -- it's not like going into another type of consciousless. You just suddenly see that YOU are the one deciding that it is important to be carrying out a particular goal; where before it seemed that it just IS important. Its like feeling like "you gotta go to work" even though you feel sick; but you don't really want to go to work; but you gotta. Then you realize that the "gotta" comes from the level of you that wants a feeling of well being and security which is achieved by having a job; but then you see that the job is just one way to achieve this perception -- and not necessarily a good one if you have to go to it feeling miserable.

This is a hypothetical example, by the way -- and probably a poor description of the experiencial change that results when you do "go up a level" -- but I'm just trying to show that it's not particularly mysterious. Its as easy as going from noticing the letters in this sentence to noticing the words to noticing that you are trying to "make sense" of them (I hope). Is your idea of the proper clinical perspective for PCT different from this concept of "going up a level"? If so, I'd like to try to understand it better; obviously, I need all the help I can get.

Best regards Rick

Date: Fri Feb 19, 1993 4:45 pm PST
Subject: predictability

[Avery.Andrews 930220.1130]

I'd suggest that the dispute between Rick Marken & Martin Taylor about the predictability of the environment might be resolved by thinking in terms of `orders' of predictability. E Coli depends on a kind of higher order, abstract predictability, namely of the fact that if moving a little bit in a given direction makes things better or worse, moving a bit more in that direction will usually produce more of the same effect. Presumably the control would deteriorate, the less often this prediction is correct. At the limiting case, where the direction you have to go to get improvement changes from microsecond to microsecond, no contol at all will be possible.

On the topic of `real' control engineers, they can certainly be very helpful (Hans Blom has already been quite helpful to me with the `technical definition of feedback issue', but it may not be the case that their mathematical sophistication will help as much as one might expect. There may simply not be much in the way of theorems that help with understanding how complex living control systems work.

Avery.Andrews@anu.edu.au

Date: Fri Feb 19, 1993 4:49 pm PST
Subject: viewpoints, internal and external

[Martin Taylor 930219 19:00] (Many Postings, including Bob Clark 930219 16:0)

On reading Bob Clark's set of levels and comparing them with Bill Powers', I am for the umpteenth time this month (especially) reminded of the great difference between the internal view and the analyst's view of a hierarchy. Maybe I am being unfair, but Bob's sound to me like the view one would see from the outside, rather than a description or model of what goes on inside an organism, whereas Bill's seem addressed to the mechanism inside the organism (again seen by an outside analyst).

(Dennis Delprato (930217)

>How is it that despite the many depictions of heating and cooling systems,
>complete with thermostat, in today's literature, virtually everyone describes
>what's going on as control of output? I suggest someone prepare a kindly
>little essay (?) spelling out how easy it is to be deceived when one looks
>from the outside in and even does a bang-up job of describing what they
>observe. Point out how one gets a very different picture when one "takes
>the viewpoint of the thermostat."

From the outside, an observer can see only actions (output). From the inside, the actor can "see" only perceptions. The two views have to be wildly different. An analyst has to pretend to be inside, while seeing only the outside of another organism.

This dichotomy of view seems to me to be at the root of the difficulty of communication with conventional psychologists, as well as at the root of many of the communication problems on CSG-L. It is SO easy to lapse into language that talks about the external view (controlling output, for example), because naturally, language is largely related to what people, communally, can agree is happening. I can agree with you that he has put a glass on the table. I cannot agree with you that he is preparing dinner, as opposed to setting the table, comparing glassmaker skills, or avoiding an accident, unless we both empathize with each other that we both empathize with "him".

We can develop social conventions about actions (including language, as Bruce's lucid 4-parter shows), but we cannot develop social conventions about behaviour (in the PCT sense). Actions are what we can see of what others are doing. We can't see what they themselves are "doing."

Bill's levels deal with different kinds of perceptual input functions. They speak, from the analysts viewpoint, about what the organism MIGHT be controlling, and have been developed by an organism that has attempted to consciously perceive what is normally unconsciously controlled. It is an empathetic view. Each level exists because there is a requirement for a different kind of perception, and the differences among the levels are (if I understand correctly) only in the Perceptual Input Functions characteristic of the different levels (I can imagine that the output functions also differ, but I don't remember that being talked about).

Bob's levels strike me as speaking to what a social contact might perceive of a person; no single ECS would act at a "Skill" level, unless I greatly misunderstand what is meant. An external observer can see skill, and the performer, **looking from another viewpoint** can assess her own skill, but no skill-level control system can be extracted from a hierarchy. Maybe Bob can describe a skill-level ECS, and prove me wrong. But I can't at the moment imagine "skill" as a level of control, in the way that I can imagine "sequence" or "program."

> As I was seeking
>to distinguish Fifth Order from Fourth Order, there was a tendency to
>consider interactions between/among individuals. Thus, with Fifth Order
>assigned to "Skills," Sixth Order could include all activities using
>combinations of Skills for purposes requiring control of interpersonal
>interactions. Examples include games, competition, cooperation,
>government, clubs, businesses, entertainment. In addition, language,
>mathematics, philosophy, systems, principles, programs are included here.

All of this is external, isn't it? You are talking about the applications for which Sixth Order systems would be used, not what Sixth Order systems do or how they are constructed. Perhaps what you are saying is that Sixth Order ECSs individually contain language models, games models, cooperation models... that they use in forming their perceptual functions. Such models are, indeed, possible. Symbolic AI depends on them. But do they belong as intrinsic components of individual ECSs?

I think I have become more sensitive recently to the importance of separating the external (analyst or observer) viewpoint from the internal viewpoint. Many of the issues raised in recent postings seem to hinge on a failure to note, and sometimes on a tendency to mix, the two viewpoints. The organism can control what it can perceive, and it cannot perceive its feedback paths, other people's perceptions or references, or its own outputs. But the analyst can perceive feedback paths and the outputs of other organisms, and can develop implausible theories that REQUIRE the organism to perceive them. S-R theory cannot work if it requires the organism to control R, for example. The analyst can see that under relatively undisturbed conditions there is a moderately consistent relationship in an experiment between S and R, as the analyst perceives them, and makes the unjustified claim that the subject produces R as a result of perceiving some transform of S. But the fact that the analyst can perceive both doesn't mean the subject can.

Many posters to CSG-L, myself included, fall into the trap of writing about something the analyst can see as if it were something the analyzed organism can see, and asserting or assuming that the analyzed organism uses that property in some way. I don't know how to avoid this problem; it is built into our language. Seeing that the problem exists is one way to avoid being caught by it. Sometimes.

Martin

Date: Fri Feb 19, 1993 9:14 pm PST
Subject: paper

[Avery.Andrews 930220.1530] Bill Powers (930218.0730)

>I think this should be preceded by a more general statement about
>the scope of PCT, lest we perpetuate the popular idea that

>control theory applies ONLY to low-level sensory-motor behavior.
>Here's a fragment for the ultimate editor of this paper to chop up:

Actually, I think that as little as possible should be seen about wider implications. Develop the right examples in the right way, and people who worth bringing over will draw the right conclusions in their own time. Emphasize the implications too much and you will (a) scare off useful but careful people who distrust big claims (b) attract an ignorant army of people whose enthusiasm outruns their understanding ((b) certainly happened to Chomsky, & he arguable didn't do enough to prevent/ameliorate it).

Avery.Andrews@anu.edu.au

Date: Sat Feb 20, 1993 12:47 pm PST
Subject: viewpoints, internal and external

[From Rick Marken (930220.1230)] Martin Taylor (930219 19:00)

Excellent post! I think that the difference between Clark's are Powers' levels may be based on more than the internal/external distinction -- but your discussion of that distinction was brilliant. I agree with you that it is probably the essense of the difference between the PCT and the conventional perspective on behavior.

Best Rick

Date: Sat Feb 20, 1993 3:21 pm PST
Subject: comments on Bruce's essay

[From Bill Powers (930220.0800)] Bruce Nevin (930218-19 A,B,C,D) --

Your position sounds pretty miserable -- my condolences. It sounds as though it's time for a change of venue.

Misery or not, your series on social conventions in language was fascinating, and as usual a model of the essayist's skill. I see it, however, as a starting point (maybe you do, too).

The reason I say this is that your review of conventions that people adopt is naturalistic, not theoretical. It is long on WHAT happens, but much shorter on HOW the WHAT comes to be, given an HPCT model. I'm not complaining about that, because I believe in the fundamental importance of the experimental-observational (as opposed to abstractly-reasoned) approach to nature and your observations certainly qualify as observations. Nor am I saying that there is an obvious theoretical approach that I see and you don't. The theory remains to be worked out in meticulous detail, a long job and one for an expert in the field.

The example of "Shibboleth" illustrates my point. How is it that the Gileadites came to say "sh" while the Ephraimites came to say "s"? According to one convention of explanation, we answer "because that is the social convention for pronuciation in the respective communities." But that is not an explanation in terms of a model; it's simply another observation, with "because" arbitrarily stuck in the middle of it. We have simply repeated the observation that Ephraimites say "s" and Gileadites say "sh."

By and large, this is the pattern of your whole essay. We observe that people who stay on the island come to pronounce vowels one way while people who leave pronounce them another way. There's no real "because" in the argument; these are observations. Either the patterns have been reported correctly or they have not; that's not a theoretical question at the level where the pattern is reported, but only a question of assessing and communicating appearances.

Similarly, when you say "Most of the tokens of membership in a social group are not accessible to self-conscious manipulation," you're making an observation that is either correct or incorrect (I assume it is correct at least in the absence of self-study). This is either true or it's not true.

When, in part B, you do begin to offer theoretical-sounding explanations, they sound like this:

>Some of the intense self-consciousness in this "age of
>perpetual embarassment" is I think due to the intimately felt
>dilemma of having to do just this--manipulate body language,
>with all the inauthenticity and inappropriateness that this
>entails--in order to create a socially viable identity or
>persona that can be a vehicle for authentically representing
>one's intentions and expectations to others, and eliciting
>cooperation from them not as a child but as an adult.

>I think this is one reason why, in these realignments of body
>language, adolescents famously tend to overshoot the mark.

Now how, theoretically, does manipulating body language in a conscious way lead to overshooting the mark? I see no reason in the antecedent statements for the claimed overshooting, unless there is some theoretical way to show that conscious manipulation of body language (as opposed to, I suppose, unconscious manipulation) is an inherently unstable control process.

If I were to accept the observations (adolescents consciously manipulate body language; adolescents overshoot the mark) and try to find a theoretical explanation within HPCT, I would do it like this (not, mind you, proposing this as a fact but only as a potentially testable hypothesis):

Body language is used as the output of a control system at a given level. It is varied as a way of having a perceived effect on other people. Adolescents have not yet developed refined perceptions of other people's reactions to them, and their body language does not often result in reactions from other people that are desired. So adolescents typically find that producing normal subtle body language outputs does not correct the perceived errors very well. As a consequence, the large errors lead to exaggerations of the body language as the child increases the output in the attempt to have the desired effect. This leads to the appearance of overshooting the mark. The system is not unstable; it is simply ineffective at controlling its inputs.

You will notice that I don't invoke any mysterious difference between an "unconsious, authentic" body language and a "consciously-manipulated, inauthentic" body language. Perhaps within my explanation you could find reasons that some usages of body language seem authentic and others not; perhaps you could invoke the association of reorganization with conscious control and anxiety to explain a certain apparent lack of skill in conscious as opposed to unconscious control (although I regard that as remaining to be established), which in turn would add to the explanation of the large errors. But this would be a basically theoretical, model-centered explanation, not one organized around informal concepts.

I would explain the illusion of hearing a glottal stop as a matter of the listener imagining pronouncing the heard sounds. If the speaker normally extends the final vowel into a diphthong, an effort would be required to prevent the final phoneme from appearing, and this might feel like a glottal stop. Listening is partly a matter of feeling how a word would be pronounced; that's why some people move their lips when they're listening to someone else speak. They're supplying the kinesthetic component of speech that they get when they speak themselves.

Your explanation in this section goes like this:

>To model this, set high gain on control for the perception of
>participation or cooperation. One form of cooperation with
>obvious utility is when an adult serves the child's needs by
>means that are beyond the child's present capacities. Another
>is when an adult guides the child through some scaled-down
>"play" version of adult skills and practices.

You can see that when you introduce HPCT, you do so only at a higher level, where the relationship of the goal to the observed behavior is several steps removed from the phenomenon. When I ask HOW this behavior is produced, I don't mean to ask WHY, but simply how, at the level of the observed phenomenon, we could account for the behavior with a control-system model.

To explain how a person imitates the sounds of language that others produce, we need go no higher than the control system that controls for sounds of a certain kind. The person hears these sounds and records them; they are then available as reference signals against which the person can compare the sounds of that person's own utterances. We might propose that a reference signal comes to be a sort of average or median of all the perceptions of the same kind that have occurred over some time span. In that case, there is little that requires further explanation: people will come to talk so they sound like most of the people with whom they most frequently associate and whose speech they most often hear. This leaves the associations to be accounted for, but that is a nonlinguistic question.

I have a general principle of explanation that doesn't often come to my awareness, but which I seem to use regularly: it's that we should look for the lowest level of explanation that actually would account for the observations. Such explanations always raise new questions, but they are of a different type -- we are always led to ask "why" as well as "how" but the "why" introduces new considerations not relevant to the basic explanation. In the example above, I explain imitation itself as a result of the operation of control systems for controlling sounds, a pretty low level. But the explanation entails the choice of people with whom one associates and talks, and that naturally requires a

higher level of explanation. However, in addressing the higher-level question, we no longer have to explain "why" people come to speak like others around them; that has been taken care of. They speak like the people they associate with, and we have a simple theoretical explanation for that phenomenon.

>I want to reach beyond the comfort zone just briefly now, with
>the observation that unconscious social convention for the sake
>of effective participation as a member of a group or community
>fits extremely well into the profile that we drew of how
>control systems of a given *order* (e.g. cells) must implement
>(and yet not experience per se) control systems of the next
>higher order.

I haven't been able to elucidate my revervations about the "superorganism" concept of society to which your line of reasoning leads. Maybe I can do a little better now.

In HPCT, the systems of a given level (which Bob Clark and I have called "order" with a different meaning) communicate with each other in a very specific way. The reference signals for a lower-level system are the higher-level system's means of action, and the perceptions of the lower level are elements from which, in part, the higher system derives its own perceptions. This is how a hierarchy of control is envisioned in HPCT.

Now just look at the relationship between the hierarchy of control systems and the neural cells of which they are composed. You see those neural cells as being a lower order of control, which operates without awareness of the functions being performed by those cells in the control hierarchy. I agree with your observation to here.

What I don't agree with is the idea that these cells are part of a control hierarchy of "orders" in the same sense that we have a functional hierarchy of "levels" in the brain. The neural signals in the brain's hierarchy do not act by setting reference signals for the individual cells, which the cells then achieve in their own perceptions, nor are the perceptions in the brain's hierarchy functions of the (chemical) perceptions inside the individual cells. So the principles of hierarchical control that we see in the brain's operation don't apply to the relationship between what you call "orders."

I'm sure that there are useful parallels between the way individual cells relate to each other and the way whole organisms relate to each other. Individual cells (or suborganizations of the same order) interact with each other as strangers; what each does to maintain its own perceptions matching its own reference signals disturbs the others through effects on the common medium. No doubt the individual cells adapt to the fact that others are present. I've had some ideas about how this can lead to specialization and differentiation.

The parallels I see, however, are those that rest on the way independent control systems interact with each other, each seeking only its own goals. All kinds of interactions can occur, not just negative feedback interactions. No cell actually sets reference signals for any other cell; all interactions are through effects on sensory inputs of other cells. This continues to be the case even when a higher order of organization appears; the higher-order organization, such as an organ, involves the evolution of a new type of cell that acts by disturbing cells that compose the lower orders, and sensing side-effects of the control operations of cells of the lower order. Cells in the pituitary gland emit hormones and sense hormones; they are a new kind of cell, not belonging to the same order as, say, skin cells.

And I believe that these orders cease with the individual organism. There is no higher-level entity, no new type of person that uses the rest of humanity for purposes unimagineable to us mortals. There is no evidence for any such thing, and even if there were, we would be quite unable to recognize it, by your own postulates.

Even in the human brain, there is a highest level. It has to exist somewhere in the brain, even if we haven't identified it yet. Beyond that level, there are no more levels of control. I think that the most reasonable proposition is that this applies to the "orders" too. And, seeing no evidence to the contrary, I propose that individual living organisms represent the highest order in the progression.

Best, Bill P.

Date: Sat Feb 20, 1993 3:28 pm PST
Subject: Tool for arm modeling" 14 df

[From Bill Powers (930220.1430)]

Att'n especially John Gardner, Avery Andrews, and Greg Williams

RE: Arm model

I now have a basic arm model with 14 degrees of freedom. It does not use dynamics or the spinal reflex model -- basically it's a massless arm set up so you can vary the 14 joint angles independently, as if each angle were under excellent servo control. The arm is plotted on the screen in 3D perspective, like in the Little Man model, and you can move the viewpoint up, down, right, and left. The limits on angular movement at all the joints are set to be realistic for a real human arm.

The degrees of freedom and the angle variables are:

```
0.  uparmpitch: upper arm elevation
1.  uparmyaw:   upper arm sideways swing
2.  uparmroll:  rotation of upper arm about axis
3.  elbow:      bending at elbow
4.  foreroll:   rotation of forearm about axis
5.  h1pitch:    bending at wrist up and down (with palm flat)
6.  hyaw:       sideways bending at wrist
7.  h2pitch:    bending of four fingers at 1st joint
8.  h3pitch:    bending of four fingers at 2nd joint
9.  h4pitch:    bending of four fingers at 3rd joint
a.  t1pitch:    bending of thumb at 1st joint toward palm
b.  t1roll:     waggling of thumb outboard and inboard
*.  t1yaw:      cocking of axis of pitch for thumb.
c.  t2pitch:    bending of thumb at 2nd joint
e.  t3pitch:    bending of thumb at 3rd joint.
```

* this angle is fixed

"d" is skipped because it designates "down" for changing viewpoint.

The program allows the user to select a joint angle from the keyboard using a number or letter from the list above, and then vary it using the x-direction of movement of a mouse. The choice "c" actually bends the last two joints of the finger simultaneously by the same amount, and "8" bends the last two joints of the thumb simultaneously by the same amount. This is just because it looks nice.

The lower-case letters u,d,r, and l move the user's viewpoint up, down, right, and left respectively.

A lower-case 'q' exits the program.

I will e-mail the Turbo C 2.0 source code to anyone who wants it. Borland ".BGI" files are required for the graphics. Otherwise the program is self-contained (no object files needed in this version). The program runs under DOS and requires a 286 or higher AT compatible -- the faster the better. A floating-point processor is highly recommended. With my 486/33, the arm moves as fast as I can move the mouse -- amazing!

The purpose of this program is to form the basis for developing more advanced control-system models of arm control. As Mary just pointed out, this model is the environment for higher-level control systems that they can act on and sense. The x-y-z coordinates of key points on the arm are continuously available for doing optical calculations.

Based on my experience with the Little Man, it should not be difficult to extend the model downward to include the actual spinal reflexes and muscle models. The physical dynamics could also be introduced -- there should be no problem stabilizing the system, as the Little Man already handles the cases that are hardest to stabilize (largest masses and interactions among moments of inertia).

The only problem is that I don't know how to write the (forward) dynamic model for the physical arm with more than 3 degrees of freedom. Also, it would be nice to be able to calculate the effects of arbitrary forces applied in arbitrary directions to arbitrary places on the arm -- another thing I don't know how to do. If any real control engineers would like to supply the missing skills I would be extremely grateful. I think we could ignore the dynamics of the hand, thumb, and fingers, at least until we want this thing to play a piano.

I have learned some things already from playing around with this model. Clearly, the main problem in adding higher levels of control to achieve coordinated movements will be in perceiving the end-result that is to be controlled, and converting errors into combinations of signals that will diminish the errors. There are obviously ways to combine control of different combinations of the joint angles to achieve relatively easy control over some regions of space, with the arm in standard configurations. As the arm passes into difficult regions, switching to controlling different functions of the joint angles should enable control to continue without requiring very complex calculations.

I can see that this model could easily become horrendously complex if we try to make it do anything useful like picking up objects. To do that, we must include not only vision but the touch feedback loop in the spinal systems, and we must put touch sensors all over the arm and hand. The spinal control loops, which are a hybrid of position and force control, look as though they will do some very interesting things when we add pressure sensors. There is negative touch feedback in all the spinal loops. When contact is made with an object, the control systems will switch automatically, I think, to a force-control mode without any specific instructions to do so. It seems likely that the fingers will have to be padded (like the real ones) so there is a small range of movement over which the pressure sensation rises smoothly upon contact.

The reason for the padding can be seen if you try to touch your upper and lower front teeth together VERY LIGHTLY. The hard contact makes the control systems unstable, and your teeth go tap-tap-tap ... . Right at the point of contact, the loop gain becomes "infinite" and the control system oscillates.

I've been reading _Robotics: control, sensing, vision, and intelligence_ by K. S. Fu, R. C. Gonzalez, and C. S. G, Lee (New YUork: McGraw-Hill, 1987). The approach seems to be to try to pack the mathematical representation of kinematics into great big four-dimensional matrices and multiply them all together to get an overall transfer function. I immediately lose any intuitive sense of what is going on -- all you can do is slog through the matrix operations and hope you did them right, which in my case is not bloody likely.

In the arm program above, I start each iteration with the arm laid out flat in a standard configuration, then apply rotations one at a time in the correct order (fingers toward shoulder) to produce a 3-D image of the arm and hand in the current configuration. This provides information about all the joint angles and all the xyz coordinates for the whole arm, in what seems to me a much simpler way. At least I can understand it.

In this robotics book there are some very general methods for computing the dynamics of arms consisting of any number of linkages. So near and yet so far -- I just can't follow what's going on, so it's useless to me, however elegant the method.

I see also that industrial robot arms have a lot more problems than human arms have. They are constructed so their joints can swivel in ways impossible for human beings, with the result of introducing lots of ambiguities that the human arm doesn't suffer. It may prove simpler to model human arms than robot arms!

At any rate, maybe some of you out there will experiment with this arm model to see what you can do with higher-level control.

Best, Bill P.

Date: Sat Feb 20, 1993 7:37 pm PST
Subject: An alternative approach to mind reading?

[From Oded Maler (930219.1830-ET)]

Apropos meanings, intent(s)ions etc., I saw this on
comp.bionet.neurosciences:

>I'm impressed by the article on page 1 of the NY Times for Tuesday, Feb.9,1993

>entitled: Computers Taking Wish as Their Command. Among the achievements
>reported in the article is the development of a computer that
>can take dictation from a person's brain by analyzing an EEG to
>determine what letters the person is thinking of.

>Any comments on my naive questions are welcome.

>Allan Adler ara@altdorf.ai.mit.edu

--Oded

Date: Sun Feb 21, 1993 1:38 pm PST
Subject: Re: Predictability and control

(Martin Taylor 930221 16:40) Rick Marken 930219.1500

We seem to be slowly converging on an understanding of the role of predictability in control. Your posting opens the door to further progress.

>Martin Taylor (930219 14:10)
>
>>Do your e-coli demonstration in an environment where the gradients
>change randomly at every time step, and the perception at any point is
>represented by a random value at every time step. Then see if it can
>control.
>
>Ok. I think what you might be saying is that there may be environments
>in which control is not possible, period. This is what you have
>described above for e coli: at every resolvable instant in time the
>input is completely random -- NO MATTER WHAT OUTPUTS ARE GENERATED.
>I don't know if there is any real environment that has these properties
>but I agree that, conceptually, nothing could live (control
>ANYTHING) in such an environment.

Quite so. I don't believe that any real environment has these properties, either. But we are considering limiting conditions here, and if I am describing a condition under which control is not possible, I am describing conditions in which life is not possible, aren't I. If you take the concept of "every instant in time" a little less rigorously, the "control impossible" situation can be equated to "too hot." When things are too hot, the molecular arrangements change to fast for control at the chemical level, on which all higher-level control is based.

More at the end of the posting, because the above is not the main issue.

>My original point was only that control does not depend on even
>partial predicability of the effect of output on input. The
>situation you describe seems less like a predictability problem
>than a bandwidth of disturbance problem -- the disturbance (or
>whatever you want to call the cause of random variation in the
>input) is occurring at a rate that exceeds the mimimum time
>resolution of the system.
>
>Rather than "a partially predictable environment" I might be able to

>go along with the constraint that control systems cannot exist in an
>environment where the only disturbance is completely random
>and of infinite frequency.

Well, that's a sufficient conditon for control to be impossible, but not a necessary one. Next, we get to the meat of the issue.

(From Rick Marken 930220.1230)

>Martin Taylor (930219 19:00)
>
>Excellent post! I think that the difference between Clark's
>are Powers' levels may be based on more than the internal/external
>distinction -- but your discussion of that distinction was
>brilliant. I agree with you that it is probably the essense
>of the difference between the PCT and the conventional per-
>spective on behavior.

OK, so we agree on the importance of considering the viewpoint from which a statement is made. In the posting to which this is a comment, I finished by saying that we all have a tendency to mix the analyst's viewpoint with the viewpoint of the analyzed organism. I think you do this in your comment on predictability.

There are two kinds of control (at least). In one kind, an ECS has by some means been connected through the outer world in such a way that an increased output has an effect on the perceptual input that is at least statistically predictable in direction. That kind of control is characteristic of the developed main hierarchy. In the second kind of control, a unitary control system I will call an RCS (Random Control System or Reorganization Control System) can determine that there is a discrepancy between a reference signal and a perceptual signal, but is not connected in such a way that the effect of a particular output on the perceptual signal is predictable. E-coli and the reorganizing system are of this kind. There is, as you say, no predictability between the output e-coli generates at any moment and the change in perception that will result from the output.

The success of an ECS in the main hierarchy depends on there being some predictability that in principle could be known within the ECS. Past outputs have had known non-random effects on its perceptual signal. An ECS controls by producing outputs that have proved in the past to move the perception in the appropriate direction, and are presumed, usually correctly, to continue to do so. As we (analysts) look higher in the hierarchy, the possibilities for predicting the effects of outputs become more complex and less reliable, but they exist. The ECSs can, in principle, model the world as it relates their outputs to their inputs. Predictability, to the ECS, is an internal matter.

In the reorganization hierarchy, all an ECS can know is that there is error, so something must be done. But what must be done? The outputs that it produces are not linked to anything that makes a predictable change in the perception. A given output may on one occasion increase the perceptual signal of an RCS, but on another occasion, the same output may decrease the perceptual signal. All it can do is to produce the output and see what happens. If matters get worse, it can produce another output quickly, but if matters get better, it may delay the next output. Eventually, with luck, it will arrive in a condition where the error is near zero, and it produces no output at all. I like to use the windblown leaf metaphor rather than the e-coli metaphor. Windblown leaves wind up in big drifts in places shadowed from winds from any direction.

The success of an RCS depends on a predictability that the analyst can see, but the RCS cannot. The windblown leaf does not know where the walls and hedges are, but the walls and hedges do not move, as a person can see. E-coli survives because it usually lives in environments that have relatively smooth and slowly changing gradients. But it cannot sense these gradients, at least not in a way that it can relate to a choice of which direction in which to move. The reorganizing system has no information that would help it to know which links to alter in the main hierarchy, but it works because (as we analysts can see) the environment in which it lives changes only slowly in relation to the (random) actions that its outputs cause. An RCS does live in a predictable environment, as an outside observer can see, but it does not overtly use that predictability, since the predictability is not related to the things it can explicitly influence by its actions. Predictability, to the RCS, is an external matter.

Let us now return to the ECS, and consider what about the environment relates to its ability to control. The environment of an ECS is its sensory inputs and its output signals (ignoring the signals that go to and come from higer levels). The ECS can know nothing about what happens between its outputs and its inputs. All it can know is that there has historically been some relationship between its outputs and changes in its perceptual signal (perhaps even its sensory inputs, but that's an unnecessary complication). This historic relation is implicit in the linkages that the ECS has with lower ECSs, as well as in the world on which it ultimately acts through those lower ECSs.

At low levels of the hierarchy, in a normal world, the predictability of the output-input relation is very high, at least in sign. Tensing a muscle will pull the attached bone, every time. But at higher levels, the effects of actions get more probabilistic in direction. The world between the output and the input is more complex, subject to more kinds of disturbance that the ECS cannot model. That's at least part of why there are different types of ECS at different levels of the hierarchy.

Now I return to your comment:

\>
\>Rather than "a partially predictable environment" I might be able to
\>go along with the constraint that control systems cannot exist in an
\>environment where the only disturbance is completely random
\>and of infinite frequency.

That's an analyst's view of the "uncontrollable" situation. In such a situation, no controller, whether ECS or RCS, could be effective. But from an organism's viewpoint, the criteria are much less stringent. All that is required to produce an "uncontrollable" situation is that any predictability in the environment be of a kind not detectable through the PIFs of the types of ECS available. We presume e-coli cannot control in the presence of rapid temporal rhythmic changes in chemical gradients. But give it a transition control system and the appropriate sensory mechanism, and it might.

As we have argued several times, the evolutionary sense of the hierarchy is to permit an increase in the stability of the organism in the face of environmental variation. Each level of the hierarchy permits control in the face of some variation in the environment that was previously unpredictable. We humans have learned to survive in a wider range of environments than any other living thing, by using our highest level control systems to develop social organizations that can make heated and insulated houses, fire-protective suits, space vehicles, submarines, and so forth. Without system-level ECSs, could we have done that?

Prediction depends on who is looking. A clever analyst may find a possibility for prediction where a simple statistical analysis does not. Random actions can find stabilities that are not directly perceived, and might even escape the analyst's first glance. But for control, predictability there must be.

And in case the point is lost again, with perfect prediction, there is no need for control. Simple outflow command will suffice.

Martin

Date: Mon Feb 22, 1993 6:28 am PST
Subject: Explanation;drafts;grades;predictions; stats

[From Bill Powers (930221.0800)] Bob Clark (930219) --

>You use two familiar, frequently used, words: "understanding"
>and "explanation." Exactly what does each "really" mean? I
>find my dictionary of little help here --

>"Explanation" seems to consist, at a minimum, of being
>classified, that is, placed in a category.

The problem with this sort of definition is that all you get is a claim that the thing to be explained is like (or at least classified with) something else, which generally is also unexplained.

>Instead of "explaining" some thing, activity, system, or
>whatnot, I prefer "description" of parts, their connections
>with each other and with other items.

I like this better. To explain a phenomenon is to describe its operation at a lower level. So models are explanations of the phenomena that they reproduce.
>The "external" aspect [of "understanding"] is more complicated,
>being displayed by asking the other party to "solve" a problem
>that requires "proper" use of the material to obtain "the"
>solution. If the result is "acceptable," it indicates (does
>not "prove") that the comparable parts of each party's systems
>are in agreement.

Yes, the question when someone says "I understand what you mean" is just what the other person's understanding is. This is the basic problem of communication.
>In my view "skill" is not a "type of control," rather it is a
>combination of perceptual variables ...

This may be a difference between our approaches that I hadn't recognized. My levels are supposed to be types of controlled perceptual variables, and by implication the systems that control them. When I label one level "programs" I don't mean just a level where programs are executed. I mean a level where we perceive WHAT PROGRAM IS BEING CARRIED OUT, and continually correct errors if we perceive a deviation from the correct program. An example would be watching people play cards. After a while, watching the play proceed, you recognize the rules in effect, and say "ah, they're playing 5-card stud." Then, if someone violates a rule of 5-card stud, you can perceive the error and (unwisely perhaps) point it out to the players to get them to conform to the rules. A rule is a form of program. To say "combination of perceptual variables" doesn't tell us much unless you say what kind of combination you're talking about.

>Thus "Skill" is a category of perceptual variables ..

I agree with that: it is a perception at the level of categories in my definitions of levels. The category level is where we use one perception (here the noise or series of marks, "skill") to refer to a collection of perceptions of lower order.

Hope the helicopter meeting is fun.

Dick Robertson (930219) --

>[Dennett] proposes that our thinking consists of "multiple
>drafts" in which different control systems (my terminology, not
>his) in the brain compete for consciousness...and therefore we
>often don't know what we think until we hear ourselves say something.

It's hard for me to imagine what the structure of a "draft" might be that it can "compete for consciousness" with another draft. How does this competition work? Do they hit each other? Bid for possession of consciousness? Grab hold of it and try to run away with it? And what does consciousness do while all this fighting over it is going on? Does it just lie there like a rag that two dogs are tugging at? Does it have any preferences for who should win? Or is it just a helpless victim? As to not knowing what we think until we hear ourselves say it, this only shows that we don't perceive output until it has an effect on input. Of course sometimes we know what we think through the imagination connection, which allows us to perceive the effect of output without actually emitting it.

>... one of my favorite speculations about consciousness - that
>it has something to do with shifting perceptions to areas in
>which errors are being corrected.

One of my favorites, too. I'd only add that if the control systems are doing well, they don't allow large errors to exist (for an appreciable-at-that-level length of time), so consciousness is attracted mainly to systems that are having problems with control and hence contain large chronic errors. By the hypothesis that consciousness directs the locus of reorganization, we have a mechanism for directing reorganization where it's needed. I presume, by the way, that you meant conscious perception -- every control system contains perceptual signals, whether they're in awareness or not.

>Are there auditory and tactile counterparts of macular vision?
>Does anyone know?

I think we find something more like the phenomenon of off-axis attention in senses other than vision. A nice demonstration is to place all five fingertips touching a tabletop, then attend to the sensations from each fingertip in turn. This is quite like attending to an object that isn't at the center of vision.

RE: Grade control

Maybe, as you say, part of the problem is that students aren't all the same. There has to be some reason for each student's presence in the class and for any studying that's done. But it seems to me that the only way to find out what it is is to interview each student individually. It seems to me that psychology tends to avoid the direct approach, preferring to sneak up on knowledge about people without letting them know what you want to know. Why not just ask? I know that subjective reports are supposed to be untrustworthy, but I don't see much evidence that other methods are any better, except maybe in psychophysics. This is sort of like saying "When all else fails, read the instructions" or "Ask the man who owns one."

>Since my self-image-control measure keeps replicating reliably
>I feel that there must be some real- time, immediate grade
>control efforts which are comparable to tracking performances
>on the one hand or self-image control on the other.

I think the difficulty here lies in individual differences. You can get an immediate test if you can deal with one person at a time, directly, as you did in the self-image project. To see what grades mean to a person, you'd have to experiment with each person -- for example, by saying "Well, I suppose I could give you a B, but your work seems worth only a C to me. What do you want me to do?" Or "Hmm, just off the bat, this score doesn't add up to anything better than a D" (when it's clearly better than that). By the person's reaction to the disturbance (if any) you can judge where the reference level is set. If a student looks too delighted or surprised at getting an A, you also learn something.

Once you get an idea of what each student's reference-grade is, if any, you can then start exploring to see how high the loop gain is -- how much effort the student would put out to correct a lower-than-wanted grade, and how much the student would relax if the grade were higher than the target. And after that you have to find out what the student knows how to do when there is an error -- does the student think there's any point in studying harder? Does other work for other classes make added studying impossible or too painful overall? Is the problem with amount of study, or with not understanding what is to be learned?

The trouble is that to do any of these explorations, you have to get to know each student very well, and to earn the student's trust so you will be told the truth. You can't keep the arm's- length relationship that's typical of psychological testing. I always had reservations about the project as you were trying to do it -- asking students to state a reference-grade doesn't tell you if the stated grade is really a goal that the person has adopted, or just something that, it is hoped, will impress you. I think your results show that the method is weak.

I think I'd begin by asking the students what grade is important to them -- would a C or D be OK, or is there really a burning desire to get higher marks? You'd have to make it perfectly clear, and be believed, that whatever the student's goal, it's truly OK with you and you won't put any pressure on to change the goal. Just to get an honest answer to this question will take a lot of effort and discussion, particularly discussion about why you're doing this and why they should trust you enough to tell you what they really think. You have to convince them that you're not going to complain if a student says "Well, basically I'm not much interested in this course, and if I can just pass it I'll have what I wanted."

Then, whatever the stated goal, the next thing is to find out what happens when they fall short of the goal or surpass it. The control system, most likely, is one-way; shortfalls are errors, but exceeding the target isn't an error. But you have to find out -- maybe if a student gets better grades than the target, the effort put in on the course will be reduced to make more time for something else -- other courses or boozing or whatever. Then it would be a two-way control system.

Finally, you'd have to find out, for those students who fall below the target, what they are doing to correct the error. Did they just lower the target? Do they know what to do? Is it possible for them to do what's necessary?

Of course to do all this you'll have to devote a lot a time to each student. Maybe you'd be better off just having them read the materials and spending ALL your time on this project. You'd just about have to, plus a lot of your own free time, plus time from assistants who understand what's going on. I don't think this is something that can be done by using ten minutes out of each classroom hour.

And finally, it's very likely that if you really carried out a research project like this, your students would change considerably. They would discover things about themselves that they didn't know. They might get better control of their time. They might find out why they have had problems with understanding the course material. You might find out that you could present it better, or in a better sequence, or in smaller but more easily- understood amounts. So by the time you got to the end of the project, your students wouldn't be the same people you started with. The characteristics of the population would have changed, maybe a lot. You would have totally messed up your experiment, and invalidated it for purposes of predicting the behavior of classes of nth-year college students with known demographic characteristics, etc.. In fact, what you say about these students would apply only to other students who had been through a similar experience, and the second time you did it you'do it better.

I guess my point is, why not just use this approach as your method of teaching? By trying to fit the approach into the standard mold of psychological experimentation, you lose the chance of helping the students with things other than the course material. There's nothing to prevent you from keeping a diary of what happens, records that you later can analyze and ponder. But when psychological testing interferes with helping students do better, hasn't the point sort of been lost?

Rick Marken (930219.0800) --

RE: predictability and control

I think the problem here is with the term "predictability." The point of having a predictable environment isn't that the control system makes any predictions:

it doesn't, or doesn't have to. The point is that if there's nothing predictable (by anyone) about the environment, then there isn't enough regularity in it to allow a systematic relationship between error and the action that corrects it.

The E. coli experiment shows that you don't have to have a whole lot of predictability/regularity. But as Martin Taylor pointed out, you do need some. For example, the sensed gradient has to be consistently positive when you're swimming toward the source of the chemical -- the judgement "I'm going the right way" has to remain valid, on the basis of what is sensed. I agree with you that most environments are at least this regular, so there isn't really a problem in most cases. To say that the regularity is required is not to imply that it's missing in any significant number of cases. This is largely an argument about possible universes, most of which don't exist.

RE: words and meanings.

There's discourse, and then there's metadiscourse. If one spends too much time in metadiscourse, discourse itself becomes impossible. Really? It really becomes impossible? How do you know that? Are you saying that "discourse" HAS AN INHERENT MEANING? Are you saying that "saying" has an inherent meaning? Hello? Is anyone there, whatever "there" means to you?

I accept as being true the idea that meanings are private experiences evoked by words. This says nothing at all about the relationship between your meanings and my meanings. A different question is how we might go about checking these meanings against each other by some means independent of the words we use to evoke them. We have all sort of ways of doing that, although the ultimate answer evades us. For ordinary discourse, I think the best approach is simply to assume that my meanings and yours coincide, but to be alert for signs that they do not. This seems to me the only way to avoid paralysis. Smart people already know that they can't just assume a veridical transfer of meaning to another person, and they already have strategies for checking this out within practical limits. I think there's plenty of checking going on on this net, although sometimes it takes a while to see where communication has gone wrong.

Martin Taylor (930219.1410) --

From Rick:
>I'm just being tenacious about this because I think
>you are trying to say that statistical properties of the
>environment place constraints on the nature (or existence) of
>control systems; or something like that.

From you:
>There's no "something like that" about it. This is exactly
>what I am trying to say.

We've never really thrashed this out. To me the environment doesn't look very statistical at all; you really have to get down on your hands and knees and look through a microscope to see any departures from smoothness. Our control systems operate at a macro level sufficiently large that statistical fluctuations just don't make much difference. A little bit, sure. But not enough so that, for example, we have to put noise into our models to make them fit behavior.

I don't doubt that there are statistical processes involved, or that your analysis of them is correct. But at the level where we presently do our modeling, I don't see the need to consider them.

I have another beef while we're on this subject. The other day, talking about focus control, you said that focus control was simple because it really involved just a one-bit change in the information. To me, this just demonstrates how the information-theoretic approach glosses over the real problems of modeling.

What you say is perfectly true, just as true as if someone noticed that the energy flow through a focus control system had to balance out to zero. But it doesn't help us understand how a focus control system would be designed. To design it, we must figure out how the "wrong" direction of movement is detected, and then how it is converted to a (one-bit) reversal in the sign of the error-to-output relationship. When we look into possible ways of detecting the wrong movement, we find that we have to perceive the current direction of output change relative to the direction of perceptual change. A function like an exclusive-or or a multiplication or a partial derivative is needed. We need a computing function that has to consider a lot more than one bit of information, even though the outcome will be only a one-bit change.

I still claim that while information theory may have true things to say about the operation of the nervous system, these are all after-the-fact characterizations, not the sort of characterizations that are of any use in designing a working system. For every theoretical limit imposed by information theory, there is a more detailed explanation of why these limits exist, an explanation that makes no use of information theory and gives us a lot more useful detail. You can characterize a filter by saying that it can pass 100 bits of information per second, max, but this doesn't tell you what the output looks like when you hit the input with a ramp, an impulse, a square wave, or a sine wave. There's no way to get from the global abstract characterizations of information theory to a prediction about the actual behavior of the system, or its actual design. To predict waveforms, you need to know the capacitance, the inductance, the resistance, the manner of interconnection, and the number of stages. And you WILL predict the waveform. What do you get from knowing only that the filter passes 100 bits per second at the most? Almost nothing, and nothing at all that you couldn't have got from the electronic analysis.

Best to all, Bill P.

Date: Mon Feb 22, 1993 10:37 am PST
Subject: how/why does language change?

[From: Bruce Nevin (Mon 930222 10:12:16)] (Bill Powers (930220.0800)) --

>The example of "Shibboleth" illustrates my point. How is it that
>the Gileadites came to say "sh" while the Ephraimites came to say "s"?

Actually, I did try to address the "how" of this. The mechanisms that I described do not match your expectations, and so go unrecognized. Perhaps I have been too sketchy. Let me try in more detail. The following is based on a summary of Labov's work on the social motivation of sound change that I wrote a few years ago.

Languages change. To account for this observable fact, people have traditionally taken a point of view looking at the end results of change. More recently, Bill Labov and his students have looked at change in progress. The two perspectives lead to different sorts of explanation. Labov's results are summarized nicely in his Sociolinguistic Patterns (1972), Chapter 7, entitled `On the Mechanism of Linguistic Change'. I suggest looking at his papers collected in that volume and in _Language in the Inner City_ to see how his results connect directly to real perceptions and perceptual control by real people in real situations, not just airy theoretobabble. Bill started his professional life as an industrial chemist. His work is solidly grounded in observation and experiment.

Viewed from the end results, sound change are changes in contrast-- splits and mergers. An example: s and sh were distinct in Hebrew (shin and samekh) and related Semitic languages. But among the Ephraimites the distinction was lost. Just so, in my wife's dialect from near Chicago, the distinction between pin and pen was lost. "Wendy from the windy city" is indistiguishable in her casual speech from "windy from the Wendy city" (though if you call attention to it, she can force a difference). The two vowels have merged for that speech community. Where my dialect has two phonemes i and e, hers has only one.

This is all that is discernible from the historical record, after a change has taken effect. But what is predominant in the study of change in progress is subphonemic variation, below the level of phonemic contrasts (what Labov calls the `cognitive function' of language elements). Differences that don't make a difference for distinguishing between words in a given dialect.

The traditional perspective views sound change as a particular outcome of "random fluctuations in the action of the articulatory apparatus, without any inherent direction, a drift of the articulatory target which has no cognitive, expressive, or social significance" (Labov, 164). On this view, a sound change (a) is followed by a period (b) of "fluctuation of forms" in which the change establishes itself, and then (c) people settle down to a new perception of regularity in the sound system disturbed only by random subphonemic fluctuations. The image of sound change arising from a bubbling pot of performance errors is perhaps analogous to the role chance mutation is thought to play in biological evolution.

Labov's findings reverse the order of (a) and (b): noncontrastive fluctuations between alternative norms--not performance errors at all, but alternative reference perceptions--give rise to sound change and thence a new regularity in the language.

As sound change proceeds in a speech community, it goes through a number of stages that Labov has identified (pp. 178-80). Stages 1-8 concern change from below the level of social awareness:

1. Undefined linguistic variable: The sound changes usually originate with a restricted subgroup of the speech community when the identity of that subgroup is weakened or threatened. Often, the form that begins to shift is a marker of regional status whose distribution within the community is irregular.

2. Indicator defined as a function of group membership: The peculiar linguistic form is generalized to all members of the subgroup, and

affects all items in a given word class, with no pattern of stylistic variation in the speech of those who use it. (The word classes involved are almost always defined phonologically, in ways relevant to the peculiarity of pronunciation involved, rather than morphologically or grammatically as in more familiar notions of what is a word class.) It takes on the status of nonverbal assertion or advertisement of one's membership in the subgroup whose peculiarity of speech it is.

3. Indicator defined as a function of group membership and age level (hypercorrection): In response to the same social factors, succeeding generations of speakers in the subgroup carry the change process beyond the model set by their parents.

4,5. Indicator for other subgroups: Other subgroups of people who have adopted the values of the original subgroup also adopt the sound change as a marker of group membership. The function of group membership is redefined in successive stages. This implies a limit of spread: the limits of the speech community, defined as a group with a common set of normative values in regard to language. (Labov's (5) is of course not a distinct stage, but rather a factor of description of a higher logical type.)

6. Marker with beginnings of stylistic variation: As the spread of the sound change reaches the limit set by the self-definition of the speech community, it becomes one of the defining norms of the speech community, and all members of the speech community perceive the significance of its use or non-use in a uniform way (often without awareness).

7. Concomitant adjustments: The change gives rise to readjustments in the distribution of other elements within phonological space. The basis for these adjustments is control for maximizing the distinctiveness of contrasting sounds.

8. Initiation of other sound changes (recycling): The concomitant structural readjustments of (7) give rise to further sound changes which in the original speech community are associated with the original change. But as other subgroups enter the speech community in (4), they adopt the original change as an established norm of the community, and treat the newer sound change (in accordance with their marginal status as newcomers) as stage (1) change. This recycling stage appears to be the primary source for the continued origination of new changes.

Stages 9-13 concern change with social awareness, and involve relations between an originating subgroup and one or more other subgroups that entered the speech community by adopting its norms in (4). Such a new group may carry the secondary sound changes of (8) beyond the level of the original change. These alternatives appear to be only partially ordered as temporal stages. I have reordered them to reflect what appear to me to be their logical interrelationships:

9. Stigmatization: If the original subgroup was not the highest-status group in the speech community, members of the highest-status group

eventually stigmatize the innovation "through their control of various institutions of the communication network." (An example is given by Labov on p. 64 n. 10.)

12. Stereotype, shibboleth: If the stigmatization of (9) is extreme, the stigmatized form may become the topic of overt social comment. As a negative model, the stereotype may be abjured, may become increasingly divorced from forms actually used in speech, and may eventually disappear.

13. Prestige Model: If the original subgroup was the highest-status group in the speech community, members of lower-status groups adopt the innovation in more careful speech in proportion to their contact with users of the prestige model, and to a lesser extent in casual speech.

10. Correction in accord with prestige model: The pattern set by the highest-status group in the speech community becomes "the pattern which speakers hear themselves using: it governs the audio-monitoring of the speech signal. . . . the motor-controlled model of casual speech competes with the audio-monitored model of more careful styles." The linguistic variable (use of the changed form vs. other forms) becomes associated with a regular stratification of styles as well as with established social strata.

    In the New York studies, it was the upper middle class (UMC) that had the highest "index of linguistic insecurity" and showed the most pronounced stylistic stratification. There continued to be no stratification in the lower class and lowest-status working class. Unmonitored, casual style A begins to become distinct for most of the working class and the lowest-status lower middle class, but full stratification of four distinct contextual styles emerges in full efflorescence only with the lower middle class. For details, see Labov, e.g. p. 128.

11. Hypercorrection in careful speech: If a subgroup that has entered the speech community by (4) lacks a form corresponding (in some word class) to a form in the prestige model, its speakers overshoot the target, pushing the sound change in that form farther than it had gone in the prestige model.

After the event, linguistics can only examine the changed structural relations of linguistic elements before and after a sound change. It can only speculate about factors below the level of linguistic contrast and below the level of social awareness. From Labov's investigations, it appears that while structural factors may provide opportunities or potentials for change, the choice for initiation and direction of change is governed by the partially opposing motivations of speech communities (a) to distinguish themselves from others and (b) to identify with other groups with higher prestige.

Date: Mon Feb 22, 1993 10:51 am PST
Subject: Re:alternative approach to mind reading?

[From Rick Marken (930222.0900)]

I'll be out of town until Friday so you'll have to figure out what to argue about all on your own.

Oded Maler (930219.1830-ET) --

The "mind reading" article is cute; are you sure it's from the NY Times and not the Enquirer? I'd be interested in 1) how they concluded that the EEG components they got were letters rather than words, sentences or other perceptions (if they were perceptions) 2) how quickly and accurately they were able to pick up the letters 3) why would the EEG component specify letters anyway? Weren't the people just imagining what they would say -- when I imagine what I want to say I don't picture the words; I "hear" the words in my head but I don't see them; and when I hear the words I don't think of what letters correspond to those sounds until it actually get's to the point where I try to type the words -- were the people in this study imagining themselves typing?

I think the description of the article shows the naivete of people's concept of behavior; the article is clearly based on the proposition that the brain (EEG waves) commmands what we do (type letters). All the snazzy physiological equipment in the world cannot overcome such a fundemental error.

I don't feel like my approach to mind reading (the test for the controlled variable) is seriously threatened.

Best

Date: Mon Feb 22, 1993 1:03 pm PST
Subject: place cells, consciousness

Bill,
I was uncertain about your meaning in your post--are you thinking that I am making up the idea of place cells or did I read you wrong? Yes, these cells do exist (?!) within at least regions of the posterior parietal lobe and the hippocampus. It's not that there is a one cell to one place correspondance. Every cell fires differentially to different regions, like cone cells fire differentially to varying frequencies of light. I see absolutely no necessity that something else needs to know which cells are firing to know place--that begs a whole bunch of questions that don't need to be asked. The firing of the place cells IS the experience of place. Nothing has to look at it. This seems obvious to me so I think I must have misundertood your point.

That no one needs to be there to watch and create consciousness is exactly Dennett's point. And I agree that it would be nice if Dennett was a control theorist cause alot of this ideas do go well with PCT, except for the MINOR (not) fact that he doesn't give much credence to volitional states in general--a behaviorist in my analysis.

I think the whole place cell way of thinking about reality really gets the idea across that we are dealing with "abstract space." The whole idea that our experience is "constructed" (I like "made up" better) is really real, not just clear. There's a "picture" (minus homuncui viewer problems in terminology) and there's the coordinates for "viewer location" which is normally centered between the eyes. But sometimes one's self (body) can be a part of the picture and sometimes the "viewer location" coordinates are not centered between the eyes--most psychedelic experiences (coordinates outside of personal space=out of body experience, coordinates slightly off="tracers" or "wavy world" or "alice in wonderland" size effects.

The latter paragraph my not relate to PCT, but its very clear to me now how the organism can reference a coordinate in space to point at or whatever. It makes the "computer coordinates" that Little Man uses ecologically valid in that we do virtually the same thing.

Well, enough rambling. Mark

Date: Mon Feb 22, 1993 1:45 pm PST
Subject: Description vs. explanation in linguistics

[From Bill Powers (930222.1230)] Bruce Nevin (930222.1012) --

I'm probably being unreasonable about this distinction between an explanation and a description, or between a model and a description. I'll relax about it pretty soon, when I'm sure that my basic point has been made (or can't be made). There's probably not much that can be done about it anyway.

>... what is predominant in the study of change in progress is
>subphonemic variation, below the level of phonemic contrasts
>(what Labov calls the `cognitive function' of language
>elements). Differences that don't make a difference for
>distinguishing between words in a given dialect.

>... noncontrastive fluctuations between alternative norms--not
>performance errors at all, but alternative reference
>perceptions--give rise to sound change and thence a new
>regularity in the language.

This is not an explanation, but an observation of a coincidence. Saying "reference perceptions" doesn't explain anything, because we still have to explain why the alternative reference perceptions arise. They are set by higher systems. What are the higher systems controlling for, that they would vary the reference-sounds as a means of doing it?

And how do "fluctuations" give rise to new regularities? I can see how they would give rise to transient new forms, but what is it that selects one of those new forms and decides to STOP the fluctuations from then on? This suggests a root explanation in terms of reorganization, but that only brings up the question of why reorganization is occurring. What sort of intrinsic error could there be such that only reorganizing the recognition and production of certain phonemes would correct it?

A further question arises as to just what the fluctuations are. You propose that they are reference-signals fluctuating. But reference signals NORMALLY fluctuate as disturbances come and go, on a short time scale. They are, after all, the outputs of higher-level systems. If a higher-level system wants to hear "Shibboleth," it turns on the system that controls the perception of that word, and alters the reference signal until the word is heard as it should be heard. Otherwise, the reference signal is set to zero.

It would seem more likely that what is changing is the organization of a perceptual function, which is more gradual and more likely to remain constant over long periods of time. What changes is not the perceptual signal -- that still matches the reference signal as usual -- but the lower-order patterns of sound-inputs that produce the SAME perceptual signal, and thus "sound right." In this way both the production and the recognition of the word would be changed, in the identical way.

If we kept trying to pin down reasonable meanings for these details, we might actually get to a level of discussion where modeling might happen.

I won't go through Labov's 13 points in detail. I am very skeptical about them from the start, because they're cast in the same old statistical population-study terms that we've been turning away from in PCT. What does it mean to say

>...it was the upper middle class (UMC) that had the highest
>"index of linguistic insecurity" and showed the most pronounced
>stylistic stratification.

?

This statement of "fact," if adopted by a modeler as a target for behavior to reproduce, would require the modeler to believe that every individual who showed high stylistic stratification was in fact a member of the "upper middle class," that everyone identified as a member of that class showed pronounced stylistic stratification (an interesting concept when applied to a model of one person), and that there were no individuals who showed the opposite relationship. The modeler would have a very hard time constructing an aspect of the model to correspond with "Upper Middle Class" in the first place. "Stylistic stratification" would be even harder, because it doesn't even apply to an individual. These are vague fuzzy terms that are essentially meaningless for purposes of modeling.

Or for that matter, for purposes of understanding human behavior. Sorry. I can't whip up a lot of interest in these statistical generalizations.

When I read things like those 13 points I have the sensation of floating in a balloon far above the landscape, with all the things actually being talked about existing way down there as little specks on the ground. And here I am, wondering how the specks work. Not much chance of figuring that out from up here.

Best, Bill P.

Date: Mon Feb 22, 1993 1:51 pm PST
Subject: AN OLD TOPIC: branching out or making a difference

******************** FROM CHUCK TUCKER 930222 ***********************

Today I just happened to be reading a book which I received a free copy of entitled SOCIAL PSYCHOLOGY IN THE '90'S by Kay Deaux, Francis C. Dane and Lawrence S. Wrightman (Brook/Cole, 1993). I was reading it to see what is the latest stuff in social psychology. Since I have given up ever using a textbook in my courses I just read them to see what is being told to most of the professors and students these days about a topic. I don't spend too much time doing it because it makes me ill. And the fact that it makes me ill is a lesson for all of us if we are interested in breaking into the mainsteam of this field or any other -- there is not one word in this book about anything that is of interest to any of us and there never will be unless we demand that it be put it there or we find some way to have these folks pay attention to what we have written. Unless we are in the textbooks that are read by thousands of professors and their students, or the HANDBOOK OF PSYCHOLOGY, or the major book in the field then no one will know what we are doing.

I love our net and would miss reading it every day but getting on another net except for just having a conversation about your work is not going to make a difference in the world of human behavior studies.

Those (I think I am talking to Rick) who want to take on the world of social psychology should write a book which translates ALL of the information mentioned in the above book into PCT terms - be my quest, I don't really have the stomach for it.

Keep writing and I'll keep reading and commenting from time to time.

Regards, Chuck

Date: Mon Feb 22, 1993 9:35 pm PST
Subject: Re: Explanation;drafts;grades;predictions; stats

[Martin Taylor 930223 0:15] Bill Powers 930221.0800

>Martin Taylor (930219.1410) --
>
>>From Rick:
>>I'm just being tenacious about this because I think
>>you are trying to say that statistical properties of the
>>environment place constraints on the nature (or existence) of
>>control systems; or something like that.
>
>>From you:
>>There's no "something like that" about it. This is exactly
>>what I am trying to say.
>
>We've never really thrashed this out. To me the environment
>doesn't look very statistical at all; you really have to get down
>on your hands and knees and look through a microscope to see any
>departures from smoothness. Our control systems operate at a
>macro level sufficiently large that statistical fluctuations just
>don't make much difference. A little bit, sure. But not enough so
>that, for example, we have to put noise into our models to make
>them fit behavior.
>
>I don't doubt that there are statistical processes involved, or
>that your analysis of them is correct. But at the level where we
>presently do our modeling, I don't see the need to consider them.
>

>I have another beef while we're on this subject. The other day,
>talking about focus control, you said that focus control was
>simple because it really involved just a one-bit change in the
>information. To me, this just demonstrates how the information-
>theoretic approach glosses over the real problems of modeling.

No, the environment doesn't look statistical. If your perceptual system normally were used to track every noise impulse, your would be using a lot of unnecessary resources. You use all the tricks available to avoid that -- one might guess initially at smoothing (don't perceive a light unless you get at least two photons from that direction in a short time), prediction -- modelling the environment, which we have talked about -- and perhaps other tricks I can't think of at the moment. The point of perception is to see what the state of the world is, not to see what the input is from the sensors. It would be odd if we normally did treat statistical anomalies as real things "out there." Why would we evolve that way? We wouldn't be able to control such perceptions, would we? But that doesn't mean we aren't working near the limits that nature imposes. The more I have thought on this (and I have been keeping, with intent to answer eventually, your postings on the matter since last June), the more I am convinced that the whole HPCT system must work to the statistical limits. If it didn;t, what would happen near zero error?

But, in respect of the one-bit focus problem that set off your response, I was only using it to show how small the lack of predictability was in that situation. It was not in the least intended to suggest HOW the control was achieved. I do agree with most of the rest of what you say. One has to get within the envelope of possibility that the information limits impose, and see just how things work. But however they work, they can't work by magic, outside the informational limits.

Martin

Date: Mon Feb 22, 1993 9:39 pm PST
Subject: Re:alternative approach to mind reading?

[Martin Taylor 930223 0:30] Rick Marken 930222.0900
>
>I think the description of the article shows the naivete of
>people's concept of behavior; the article is clearly based on
>the proposition that the brain (EEG waves) commmands what we
>do (type letters). All the snazzy physiological equipment in the
>world cannot overcome such a fundemental error.

Well, many years ago, I did see a demonstration at McGill University in which the subject "thought" a square on a screen (a cursor, if you like) into one of the four corners of the screen, using the EEG. It worked, and to say on purely theoretical grounds that it didn't is a rather fundamental error. I know nothing of more recent work in the area, but it doesn't seem outrageous to me that people could "think" words onto a screen in the same way, if the vocabulary were small enough.

Martin

Date: Mon Feb 22, 1993 10:01 pm PST
Subject: Re: Description vs. explanation in linguistics

[Martin Taylor 930223 0:45] Bill Powers 930222.1230
>
>I won't go through Labov's 13 points in detail. I am very
>skeptical about them from the start, because they're cast in the
>same old statistical population-study terms that we've been
>turning away from in PCT. What does it mean to say
>
>>...it was the upper middle class (UMC) that had the highest
>>"index of linguistic insecurity" and showed the most pronounced
>>stylistic stratification.
>
>?
>
>This statement of "fact," if adopted by a modeler as a target for
>behavior to reproduce, would require the modeler to believe that
>every individual who showed high stylistic stratification was in
>fact a member of the "upper middle class," that everyone
>identified as a member of that class showed pronounced stylistic
>stratification (an interesting concept when applied to a model of
>one person), and that there were no individuals who showed the
>opposite relationship.

Bill, I hate to say it, but I think you are missing a key point here. This is a quite different aspect of statistics from what "we've been turning away from in PCT." What is happening in this situation has nothing to do with whether any individual qualifies within the class and showed the relationship. It has to do with what individuals are exposed to, their perceptual environment. If someone associates with people, many of whom act (not behave) in a certain way, that individual will experience less conflict if perception of those actions conforms to that individual's reference levels for that kind of action. If the individual has those kinds of reference levels, the individual will tend to emit similar kinds of actions (on a statistical basis, actions being situation-dependent). Others exposed to that individual will experience less conflict if they, too, have like reference levels. Less conflict, less reorganization. It is a statistical question about who individuals come in contact with that determines the likelihood that the individual's reference levels will lead to conflict and thus reorganization. The net result is that very small statistical fluctualtions in the action probabilities of the contact group can lead to dramatic changes over time in any kind of conventionalized activity, including language.

The development of conventionalized forms seems to have an element of positive feedback about it, in which small fluctuations can grow into defining characteristics of a sub-group. Obviously this can't go on forever, so there is some limiting condition. It would be interesting to try a simulation of mutual reorganization some day, to see whether anything has to be added to the straightforward view of reorganization to accommodate this kind of social development. (I don't think it will be me that does it!)

So, while we can keep the main discussion of the INTERNAL effect of statistical fluctuations in abeyance, I think in this case the importance of statistics can be argued directly.

Martin

Date: Tue Feb 23, 1993 12:03 am PST
Subject: Re: meaning, control, prediction

[From Oded Maler (930222.1900-ET)] Rick Marken (930219.0800)

* I don't get it? Are you saying that this statement, in particular,
* reflects a belief on my part that words themselves have meanings? Is it
* the word "fact" that causes a problem? Do you imagine that the
* "evoked meaning" view of words rules out the idea that there
* are "facts"? I didn't mean to imply "ultimate reality" here. The
* fact that control is the control of perception is a fact in then
* sense that it is an experience that you can reliably produce for
* yourself in various ways -- through demonstrations (like those
* described in "Mind Readings" and Bill Powers' Demo program) and
* mathematical analysis. It is a much a fact as the "fact" that
* electric current flows from a point of high to a point of low voltage,
* the rate being proportional to the voltage difference and inversely
* proportional to the resistance of the medium.

But sometimes you seem to imply that the evoked meaning of the words B, C, and P within yourself have a privilleged status "closer" to their ultimate meaning, while for other people the evoked experience of those words is different - e.g., for many people any trajectory of a dynamical systems is a "behavior". You might claim, correctly maybe, that what the words evokes for you is more interesting or relevant for the study of living systems, but as you sometimes put it, it may look to people who don't share your experience as if you are claiming to have direct access to the ultimate meaning of such words.

As for the analogy with currents and voltages, I'm not sure how far you can push it. The latter speak (approximately) on what you can measure with some instruments in certain locations on the circuit. But I'm ready to be enlightened.

[Bruce Nevin]

Thanks for the quote.

Incidently I came across IEEE engineering in medicine and biology magazine, december 92, and the issue is dedicated to posture control. In addition there is an announcement of a conference on biomechanics and neural control of movements, Ventura CA, 16-31 July 92. It contains extended abstracts of most of the talks that will be given there by various mainstream robotic/motor control researchers. The titles include "Control of contact in robots and biological systems" (Hogan), Robotic vs. Human movemment Control (Inbar) and others. One of the abstracts (Prochazka) promises to discuss (among other questions) the question: "when is a predictive open-loop control superior to complex feed-back control?"

--Oded

Date: Tue Feb 23, 1993 4:43 am PST
Subject: ???

I received this:

>Bill,
>I was uncertain about your meaning in your post--are you thinking that
>I am making up the idea of place cells or did I read you wrong? Yes,
>these cells do exist (?!) within at least regions of the posterior
>parietal lobe and the hippocampus. It's not that there is a one cell

I assume you mean another Bill. Please use last names.

Bill Silvert

Date: Tue Feb 23, 1993 5:49 am PST
Subject: Naming names

From Bill Cunningham 920223.0835 Bill Silvert--

You are right, of course, about specifying the Bills. But take cheer. There is another Bill Cunningham working at Ft. Monroe, whose userid is `cunningw'. Wanna compare misaddressed mail? The solution is simple: all payments and accolades go to me; all dunning notices, complaints and nasty tasks go to my Doppelganger. Of course, he uses same strategy.

Bill C.

Date: Tue Feb 23, 1993 9:16 am PST
Subject: Place cells;consciousness;statistics;language usages

[From Bill Powers (930223.0800)] Mark Olson (930222) --

RE: place cells.

>I was uncertain about your meaning in your post--are you
>thinking that I am making up the idea of place cells or did I read you wrong?

I believe you; I'm just questioning whether the existence of place cells is sufficient to account for all the phenomena involving place. Consider the simple problem of deciding which of two places is farther away from you. Somehow the signal from one place cell must be compared with the signal from another place cell (by a higher-level system) to provide, at the very least, an indication of "this closer than that." So there must be some way for a higher-level perceptual function to perceive this relationship. If all that the signal from a place cell does is to indicate "something here", there is no way for a higher system to get quantitative distance information directly from the signals. And if such distance comparisons can be made, then there is a higher system that can reduce a set of "place" signals to coordinates and represent the value of each coordinate by a single signal. We have then really bumped place perception up a level.

>It's not that there is a one cell to one place correspondance.

>Every cell fires differentially to different regions, like cone
>cells fire differentially to varying frequencies of light.

This makes it sound even more as through the placeness of the signals (in the brain) is irrelevant -- that actual place perception (location relative to a background coordinate system) happens at a higher level. It may be that certain computations involving geometry require a lot of parallel processing and that to carry out the computation it is best to have the computing elements as close as possible to others with which they interact. This would lead to the appearance of mapping from external scenes to layers of the cortex, without that mapping itself having any geometrical significance (the cells could be geometrically scrambled without altering the computations, but then all the pathways would be longer than minimum length).

I don't mean to minimize the importance of the fact that positions within images do map onto positions within neural nets, more or less. But we have to realize that a "place cell" can't itself be computing placeness. A place cell is merely a location where you can measure a signal that corresponds to a function of geometical position. This cell represents the output of an underlying computing network of cells; it is the underlying computing network, not the cell that receives the outcome of the computations, that is computing the value of a place variable. You wouldn't call the loudspeaker of a radio a "music cell" just because that's where the music comes out. I think that Heubel and Weisel's designation of certain cells as "responding to orientation" was unfortunate. Clearly, they're simply the cells that receive the output of orientation calculations.

Perhaps, at the lower levels in the brain, lower-level information like intensity and color are being extracted, and that the mapping that goes on from level to level is simply preserving geometrical information so it can reach higher systems that reduce spatial relations to single signals. We know, for example, that retinal cells are "place cells" par excellence. The signals from retinal cells map onto layers in the midbrain where, if I remember correctly, attributes such as objectness are extracted and represented as signals. It isn't until you go one more level that spatial relationships, directions of motion, sequences, and so on are computed.

Even as you go higher in the visual systems, into the visual cortex, mapping seems to be preserved. But this mapping may represent the fact that lower-level attributes of the visual field must remain associated with positions in the visual field so you can perceive, at the same time, a "blue triangle" "next to" a "red circle." In other words, the mapping doesn't itself provide the perception, but simply preserves information until the signals reach the level where the proper many-to-one computation can be performed.

>That no one needs to be there to watch and create consciousness
>is exactly Dennett's point.

It's not as simple as Dennett and others make out. The mere presence of a signal in a perceptual pathway is not enough to assure that the signal is in awareness. The simplest example is perception (and control) of breathing, which until I mentioned it was not in your awareness. But the perceptual signals must have been there all the time, because control of breathing was going on.

This is even more evident when you consider levels of perception. Right now you're reading these printed marks, but even though you recognize the words, you aren't aware of all the lower-level details of the perceptions such as the height or brightness of the letters. Clearly, however, if your nervous system were't continually sorting out such lower-level details, there would be no perception of the words, and without the words you wouldn't be getting the meanings right now. If your kinesthetic control systems weren't controlling the signals representing muscle tension, joint angle, and so forth, you wouldn't even be able to sit up to read the words. You CAN become aware of these signals, but awareness is not required for your brain's control systems to be controlling them.

Phenomena such as these tell me that awareness is something different from the existence of perceptual signals. It can move around in the brain. Consciousness, as I think of it, has two components. One component is the set of perceptual signals that are functions of lower signals and so forth to basic sensory information. The other component is awareness. Only when you put these two things together do you get what we call consciousness, and then the CONTENT of consciousness is set by the neural systems involved in awareness. Other neural systems not so involved go right on working -- but the perceptual signals being controlled are not in awareness.

Chuck Tucker (930222) --

>Those (I think I am talking to Rick) who want to take on the
>world of social psychology should write a book which translates
>ALL of the information mentioned in the above book into PCT terms - be my guest

This is much the way I feel about it. The real revolution has to be carried out by young enthusiasts who are willing to fight the dragon. When I read books like the one you mention, I just get tired. What we need are nasty young people who will stand up in a class and say "I don't see any reason to believe that a word you're saying is true." This, of course, will not do their careers or reputations any good, but as Arlo Guthrie said (after Gary Cziko's observation), if enough of them do it the definition of good psychology will have to change.

Martin Taylor (930223.0015) --

>No, the environment doesn't look statistical. If your
>perceptual system normally were used to track every noise
>impulse, your would be using a lot of unnecessary resources.
>You use all the tricks available to avoid that -- one might
>guess initially at smoothing (don't perceive a light unless you
>get at least two photons from that direction in a short time),
>prediction -- modelling the environment, which we have talked
>about --

This almost defines the proper level at which to do modeling of behavioral organization. Start with muscles: the muscles themselves average over statistical ensembles of driving impulses, summing them up into a single applied force. This tells you to model the spinal control loops on that same time-space scale, ignoring individual impulses and the fact that a "single" control system is really a collection of parallel control systems. The perceptual signals representing muscle force are really strings of rather sparsely-space impulses in individual channels -- but average them over 50 milliseconds and over 100 or so parallel channels, and you have a nice smooth perceptual signal with a wide dynamic range.

As you and I agreed a year or more ago, the brain's neural systems don't actually do any statistical calculations. They work in terms of signal averages (the neurochemical nature of neuron operation helps to do the averaging). It's just that the operation of physical systems can handle impulses and filter them through analogue neurochemical processes, and that the resulting input-output relations correspond to certain statistical measures analyzed with symbolic calculations. The nervous system is basically a physical system; it operates because of direct physical and chemical interactions, not because of the symbolic ways in which we represent, approximate, and analyze those interactions.

There is a level of analysis of neural systems at which we speak of continuous signals, not discrete impulses, and single channels, not parallel redundant channels. At this level of description, we see a continuous system with a low noise level. There's a sort of complementarity going on here: we know that at a more detailed level of description, on which we can resolve sub-millisecond changes, the whole system is a roaring torrent of noise. However, what we see on this time-scale has essentially no relationship to behavior as we normally observe it.

We can also go to a more abstract level of analysis, in which we ignore the details of behavior and look only for global patterns and generalizations. This level, too, fails to represent behavior as we normally see it -- one person at a time.

Somewhere between the most general possible description of behavior and the most detailed, there is the level at which I try to work. It's the level of the circuit diagram, not the level of holes and electrons; the level of the computing component, not the level of the whole organism. I think this level of description is typical of engineering. The problem is to represent the whole system as a collection of functions, where each function is well-defined and behaves regularly. By characterizing those functions clearly we can assemble them into behaving systems without worrying about their interior details. When the right level of analysis is found, the system becomes understandable in terms of functions we can understand because they are simple.

If we go too far below this level of description, we begin to drown in detail; too far above it, and we "say less and less about more and more."

>... in respect of the one-bit focus problem that set off your
>response, I was only using it to show how small the lack of
>predictability was in that situation.

A one-bit error does not necessarily imply a small prediction error. How small the error is depends on which bit is in error. If you are wrong about the most significant bit (which is what the sign bit is), the error can be the largest error possible. That's what happens when you get the sign of the output function wrong in a control system.

RE: Labov's 13 points

>What is happening in this situation has nothing to do with

>whether any individual qualifies within the class and showed
>the relationship. It has to do with what individuals are
>exposed to, their perceptual environment. If someone
>associates with people, many of whom act (not behave) in a
>certain way, that individual will experience less conflict if
>perception of those actions conforms to that individual's
>reference levels for that kind of action. If the individual
>has those kinds of reference levels, the individual will tend
>to emit similar kinds of actions (on a statistical basis,
>actions being situation-dependent).

I think I agree strongly. What you're doing here is taking the viewpoint of the individual, not that of the externalized observer looking at the whole group. To the individual, "the group" is all of the others, with the boundaries set by the way the individual perceives groupness. We can understand what "group" means only by understanding how the individual perceives the group.

Once that is settled, then what you say about the statistical basis of the relationships follows (although I put this in terms of perceiving average situations, the equivalent of a statistical treatment). It is the individual "doing statistics" on interacting with all the others that reveals regularities, perceptions of the group. It isn't likely that the individual does this by the same means or under the same assumptions as the statistician would (Gaussian distributions, covariances, least squares, multiple factor analysis etc.), but perhaps what the statistician might deduce has some resemblance to what the individual might deduce.

Having established this much to our satisfaction, we can then start down the road toward a model. Just what perceptions of the group DOES the individual control for? How do these perceptions relate to non-social reference signals (like getting enough to eat)? What are the potential conflicts that would lead an individual to adopt, or avoid adopting, certain of the mannerisms of others? By trying to put together a coherent picture at this level of analysis, we might actually come to understand how and why the individual comes to speak like, or differently from, the surrounding groups.

Even reading the words of a linguist -- Labov -- we can see that the answers to these questions come from territory that includes far more than language.

>The development of conventionalized forms seems to have an
>element of positive feedback about it, in which small
>fluctuations can grow into defining characteristics of a sub-group.

Well, it COULD. If the situation improves because of adopting some different form and that improvement leads to adopting still more different forms, there will be a positive-feedback flip to a new set of forms. But if an improvement in the situation leads to less motive for adopting new forms, there will be no positive feedback. Almost by definition, the "flip to the new set of forms" can't persist unless the result is to restore negative feedback.

This, by the way, illustrates my point that social systems are not hierarchies of control systems. All relationships, including positive feedback, are equally possible in a social system.

Oded Maler (930222.1900-ET) --

Did I not send my post relating to metadiscourse? Nobody's commented on it yet. My point was that we have to use metadiscourse sparingly if we want to avoid paralysis. Meanings are evoked by communication, not transmitted, sure. But we have to assume some commonality of the meanings in order to talk at all. We spend a lot of time trying to check up on the meanings others get from our words. Most of the time we're satisfied with the result.

Bill Silvert (930223) --

Bill, my understanding of our conventions (as first suggested by Gary Cziko) is that we preface replies to a specific person by mentioning the full name and time of transmission (as in the subheader just above) the first time it occurs, and can then go back to first-name references until another person with the same first name is mentioned. In the post from Mark Olson, this convention was not used, and so led to your confusion.

There's another convention that I've found very convenient and to which I try to stick. At the beginning of a post, in the very first line, the author of the post says who the post is from. I use the convention

[From Bill Powers (930223.0800)]

This allows people who save transmissions to edit out all the header material that the net puts in, to make more compact files. When you're reading a long list of posts concatenated together (as I do), it also lets you know who's talking without having to page down to the end of the post to find the signature (assuming the sender remember to put one there). In a short post where the signature is visible on the same screen as the beginning, this may not be necessary.

Best to all, Bill P.

Date: Tue Feb 23, 1993 9:54 am PST
Subject: Re: Place cells;consciousness;statistics;language usages

[From Oded Maler (930223.1815-ET)] Bill P (930223.0800) Oded (930222.1900-ET)
* Did I not send my post relating to metadiscourse? Nobody's
* commented on it yet. My point was that we have to use
* metadiscourse sparingly if we want to avoid paralysis. Meanings
* are evoked by communication, not transmitted, sure. But we have
* to assume some commonality of the meanings in order to talk at
* all. We spend a lot of time trying to check up on the meanings
* others get from our words. Most of the time we're satisfied with
* the result.

Feed-back along slow communication links is not so reliable... I recieved your post today, but I answered Rick only yesterday. I agree with you (and of course practically there is no other way) but unfortunately the meta-level is my favorite level...

This may motivate some search for something absolute independent of the collective perception of some arbitrary bunch of primates. But let's not get into that.

Apropos the biomechanical references I mentioned yesterday, ther is another one there, by A. Prochazke, "Task-dependent scheduling of feed-back gains ia animal motor systems: Too much too soon for FES control? (FES is functional electrical stimulation, as a method for restoring muscle activity in paralyzed patients).

The words are different but the evoked experience seems close to some beginner's version of PCT. "some parts of the brain might be specialized for setting transmission parameters elsewhere".

I'll go tomorrow to Germany and on monday to .. Paris. Hope out file system will not crash under the pressure of the CSG mailbox..

--Oded

Date: Tue Feb 23, 1993 9:59 am PST
Subject: Moyer program Mind-Body Connection on PBS

******************* FROM CHUCK TUCKER 930223 ************************

Last nite on PBS Moyer had a wonderful program which showed that the old dualism of mind and body hampers the understanding of human action. There were so many times while I was watching that I said to the TV "You fool don't you understand PCT - that solves the problem!!" Moyer will have another program tonite and another on Thursday (they are on at 9 PM here). I plan to watch them and would recommend them on the basis of the first one hoping that the others are as informative. There is also a tape and book that comes with the programs. YOU MAY NOT AGREE WITH ALL OF IT BUT I THINK IT IS WORTH A GANDER.

Regards, Chuck

Date: Tue Feb 23, 1993 11:00 am PST
Subject: language norms

(Bill Powers (930220.0800) (930222.1230)) --

Sorry to seem obtuse. I'm trying to give you the data as they are available. I intend to translate more than I in fact do. It's not just laziness. I shouldn't be doing this: I should be doing things more directly ensuring my survival here.

Dialect: a manner of pronouncing, etc., that is characteristic of an idenfied subpopulation. To a child coming up in that subpopulation, perceptions of how "my people" talk; thence, reference perceptions that determine how I talk.

Linguistic insecurity: a measure determined by discrepancies between how I say/believe I talk and how I actually do talk when off my guard. People who are members of upwardly mobile/downwardly vulnerable social strata characteristically evince more linguistic insecurity. In formal situations, or when attention is called to their way of talking for any reason, they emulate ways of talking that were the careless norm of the previous generation of the social strata above them; let the formal interview, etc., be interrupted by a telephone call from a peer, etc., and their way of talking is controlled for conformity to the less-esteemed norms of their social stratum. I say that I say "get"; I hear myself say "git", with some embarrassment. People with this "linguistic insecurity" maintain reference perceptions for at least two dialects. They appear to control their speech with respect to one under conditions of embarrassed self-watchfulness , emulating esteemed models; with respect to the other under conditions of cooperative co-membership with peers.

Once children acquire their reference perceptions for control of dialect, they do not refer to the speech of those around them for further guidance. Adults thus have as internalized norms for their manner of talking perceptions based upon the speech of an older generation, extant when they were children. These norms appear to generalize from and idealize actually heard characteristics of speech. It is in this generalizing and idealizing that hypercorrection and other "overshootings of the mark" take place. And I said based upon the speech of an older generation, yet too children calibrate their new norms to one another. I can only guess at the mix of factors; no one has done the basic observational work, to my knowledge. It is clear that when adolescents on Martha's Vineyard came to choose what kind of person to be, they already had in memory, ready to use, reference perceptions of how each kind of person talks, not as acoustic images of specific words, but as perceptions to control this way for this dialect (open [a] in words like "house") or that way for the other (close [A] ("uh") in those kinds of syllables). That the controlled perception is applied to a broader class of syllables goes unnoticed (in "knife" and "wide"), so we can't be dealing with a repertoire of acoustic images of words. The generalization and idealization applies across the board to all words containing these diphthongs--however that is perceived and the perception controlled.

This process of constructing new norms only is possible at certain junctures in life. Moving a Texan to New England may well eradicate the Texas accent in a child, or an adolescent (assuming she or he embraces membership in the new community of peers), but almost certainly not in an adult.

>I would explain the illusion of hearing a glottal stop as a
>matter of the listener imagining pronouncing the heard sounds. If
>the speaker normally extends the final vowel into a diphthong, an
>effort would be required to prevent the final phoneme from
>appearing, and this might feel like a glottal stop. Listening is
>partly a matter of feeling how a word would be pronounced;

Of course that is so. But "the speaker normally extends the final vowel into a diphthong" and can imagine stopping short only by means of a final consonant if her native language is English; if her native language is e.g. French, or Russian, the speaker does NOT (equally "normally") extend the final vowel into a diphthong, and can easily imagine simply stopping a word-final accented short vowel, no consonant needed or expected. The point that you are missing is, what is normal? Sapir shows how it is related to patterning observable in the language. It appears that we learn patterns in what we observe and use those patterns to predict (and guide our creation of) utterances far beyond what we have observed. As children, we generalize and idealize from what we observe to patterns that are different in detail from those arrived at by our parents' generation when they were children. We observe what they do and what they say they do and whom they emulate and fail to emulate and when.

I'm not putting this very well or convincingly, I'm afraid. I can't. I have to get back to packing before the day runs out on me. I'll have to trust your forbearance to look at the messy facts and participate in making PCT sense of them, rather than taking me to task for failure already to have done so for you.

Gotta go, Bruce bn@bbn.com

Date: Tue Feb 23, 1993 11:53 am PST
Subject: Re: statistics

[Martin Taylor 930223 14:20] Bill Powers 930223.0800)

I don't think we disagree very much on the statistical issues, but the little disagreement I do perceive may mask a larger one, so I'm going to wiggle the niggle (or as you once put it, "pick the scab").

>Somewhere between the most general possible description of
>behavior and the most detailed, there is the level at which I try
>to work. It's the level of the circuit diagram, not the level of
>holes and electrons; the level of the computing component, not
>the level of the whole organism. I think this level of
>description is typical of engineering.

This I agree with. But the problem is one of determining what "circuits" are likely to be in a living organism, not one of designing circuits to perform a task. The method is to design the circuits to perform the task and then see whether the performance is like that of the living organism. Most of the time, that's a question of seeing where and to what degree both fail to perform the task correctly, since there are in most cases many ways of performing the task correctly.

The limiting factor in the operation of a control system is its precision of measurement (and possibly of execution, though most of that can be eliminated by the feedback mechanisms). Precision involves time and the gathering of information about the thing being perceived. You can't determine that there is a positive or negative error if the error is small and the perceptual precision poor. I quote:

> Start with muscles: the muscles
>themselves average over statistical ensembles of driving
>impulses, summing them up into a single applied force. This tells
>you to model the spinal control loops on that same time-space
>scale, ignoring individual impulses and the fact that a "single"
>control system is really a collection of parallel control
>systems. The perceptual signals representing muscle force are
>really strings of rather sparsely-space impulses in individual
>channels -- but average them over 50 milliseconds and over 100 or
>so parallel channels, and you have a nice smooth perceptual
>signal with a wide dynamic range.

This is quite right, and it illustrates my point. The limiting factor in the dynamic performance of the control system is the statistical noise--the information rate. We do consciously perceive the smooth, apparently noise-free percepts. In fact, some experiments seem to show that we do it "unjustifiably" well (the phenomenon is called "perceptual sharpening"--we see what isn't there in the data, using imagination to fill in the gaps, and we do this when the informtion in the sensory signal is enough only for a moderately probable guess at what to imagine).

When a control system is trying to keep a percept near a reference, it can do so only to the extent that it can imagine or analyze the incoming sensory data. The limit of precision is seen as "precise." When I don't wear my glasses, things don't look fuzzy, but I can't read small print. I know that if I put my glasses on, I will be able to. Things don't look more precise, but smaller things look precise.

>As you and I agreed a year or more ago, the brain's neural
>systems don't actually do any statistical calculations.

I still agree with that.

Does this open or close the scabby wound?

Martin

Date: Tue Feb 23, 1993 1:30 pm PST
Subject: Language; designing robots vs. analyzing organisms

[From Bill Powers (930222.1300)] Bruce Nevin (930222) --

>Sorry to seem obtuse.

You? Never. I, on the other hand, can chew over a point past where it's recognizeable.

In an earlier post I brushed past a point I'd forgotten to make. It's that people speak as they do for reasons that include nonlinguistic reasons. You bring them up yourself:

>Linguistic insecurity: a measure determined by discrepancies
>between how I say/believe I talk and how I actually do talk
>when off my guard. People who are members of upwardly
>mobile/downwardly vulnerable social strata characteristically
>evince more linguistic insecurity.

At one level, we want to explain how people can come to speak in a certain way. That's not very hard to explain with PCT, I think. They make the sounds they hear match the sounds they intend to hear.

But at another level, we have to ask "Why do they want to hear one set of sounds rather than another?" The answer to this question can't given on linguistic grounds, but on the basis of relationships with other people, self-image, and so forth. That's what you do in the quote above. "Insecurity" isn't a linguistic concept, nor are such concepts as upwardly/downwardly mobile or social strata. That's most of what I'm trying to get at. We have to model the perception and control of social interactions, with language as one facet of them but with many other considerations of equal importance. Why should one be embarrassed at saying "git?" This has nothing to do with language per se; you'd be at least as embarrassed at finding yourself picking your nose in the middle of delivering a speech. The norms we adopt pertain to all sorts of behavior, and our perception of "likeness" to those whom we want to be like include far more than the way they talk. The GENERAL problem we want to solve has to do with the perception of likeness and where people set their reference levels for it, not any particular example of it.

>... "the speaker normally extends the final vowel into a
>diphthong" and can imagine stopping short only by means of a
>final consonant if her native language is English; if her
>native language is e.g. French, or Russian, the speaker does
>NOT (equally "normally") extend the final vowel into a
>diphthong, and can easily imagine simply stopping a word-final
>accented short vowel, no consonant needed or expected.

That what what I intended to communicate -- that the illusion of the glottal stop would appear only in speakers who had to prevent themselves from extending the vowel into a diphthong in order to imagine matching the heard vowel.

>The point that you are missing is, what is normal?

I don't think I'm missing it, unless I'm missing it in some subtle way. "Normal" is simply how you hear/speak. You, the speaker of a particular language.

>It appears that we learn patterns in what we observe and use
>those patterns to predict (and guide our creation of)
>utterances far beyond what we have observed. As children, we
>generalize and idealize from what we observe to patterns
>that are different in detail from those arrived at by our
>parents' generation when they were children. We observe
>what they do and what they say they do and whom they emulate
>and fail to emulate and when.

So we invent patterns as a way of putting order into what we hear, right? Once having invented a pattern, we apply it beyond the original scope in which we invented it. Each person does this inventing a little differently, but with so many people around it's likely that many inventions will be similar (particularly considering the similar mechanical constriants on actually producing speech). And there's the phenomenon of hearing the way other people pattern their speech, which provides potential reference signals. If we could model how ONE person does this, we would then be able to put several such models together, interacting, to see what equilibrium conditions would emerge.

I feel that before we solve this problem, a lot of time will pass. All we can do is keep putting one foot after another and try to go in the right direction.

P.S. An afterthought on "Dialect: a manner of pronouncing that is characteristic of an identified subpopulation."

When you eliminate the dialect itself as a means of identifying the subpopulation, what is left? Location of residence, income, sex, age, education, occupation -- all the marks that traditional psychology has tried to use as an "objective" way of identifying populations. It is notoriously difficult even to define a "population" in psychological experiments -- why should it be any easier for linguists? And why should these marks of membership in populations work for linguistics as predictors of verbal behavior any better than they do in psychology as predictors of any kind of behavior? I look askance at all purported "facts" that were established by these traditional means.

Martin Taylor (930223.1420) --

>... the problem is one of determining what "circuits"
>are likely to be in a living organism, not one of designing
>circuits to perform a task.

Yes, and this is a major difference between practical robotics and modeling of human behavior.

There's a grey area, however, in which we don't have any knowledge of how a human being OR a robot might accomplish a particular kind of behavior. In that case there are great advantages in knowing how to design artificial circuits. When you design a circuit that will accomplish the behavior, you know at least ONE way it could be done. If you have a lot of experience with design, and can't think of ANY way a particular behavior might be produced, this can be a hint that perhaps you're characterizing the behavior wrongly.

I think this is my basic reason for rejecting the "inverse- dynamics" school of motor programming. When I think about actually designing such a system (assuming lots of expert advice to get over mathematical hurdles and plenty of money for building models), I start laying out the main features of what would be necessary. The system would have to have information about mechanical properties of the arm like mass distribution, joint friction, and so forth. It would have to sense the biochemical state of the energy supply and the muscles, and be able to compute the effects on muscle responses to driving signals. Loads would have to be sensed, and the proper transformations applied to compute compensations. Disturbances would have to be predicted on the basis of quantitative observations of the environment and complete knowledge of the laws of physics and chemistry involved.

In fact, the more I think about actually designing such a system that would behave realistically in the real world, the more clearly I realize that I couldn't do it with ANY amount of expert advice. I couldn't do it even if the expert advice were correct, because the physical components with which I'd have to build the system wouldn't have the necessary precision and wouldn't be able to detect all the required information.

So as an engineer, I simply say "That can't be how it works." I could say this even without having an alternative proposal, simply because the attempt to design an artificial system with these properties would clearly fail. The more you know about designing real systems, the less plausible the inverse-dynamics solution appears as a model of a living system.

>The limiting factor in the operation of a control system is its
>precision of measurement (and possibly of execution, though
>most of that can be eliminated by the feedback mechanisms).
>Precision involves time and the gathering of information about
>the thing being perceived. You can't determine that there is a
>positive or negative error if the error is small and the
>perceptual precision poor.

Complete agreement. In fact in this paragraph you pretty much state the way we should evaluate living control systems. There are artificial systems that perceive orders of magnitude more precisely, and actuators that are equally superior to human ones, and methods of signal transmission with 50 db greater dynamic range. We have to define our scales of measurement to reflect the system we deal with. An error of 2% is "small"; an effort of 100 Kg at the point of application is "large." Perceptual information that is accurate to 5% and has 2% noise in it is good information. We should expect normal control behavior to reflect these limits, and only worry about noise and imprecision when they are significantly greater. A world that we apprehend within these limits of precision and noise appears smooth and regular to us. That's the world I'm trying to model.

So I completely agree when you say that the limit of precision is seen as "precise." If we both keep this in mind we should be able to keep the domains of analogue modeling and statistical analysis properly separated.

Best to all,   Bill P.

Date: Tue Feb 23, 1993 3:42 pm PST
From: CZIKO Gary   MBX: g-cziko@uiuc.edu
TO: * Dag Forssell / MCI ID: 474-2580
Subject: Senge Program

Dag (direct):

This Friday a satellite TV broadcast will be shown here called "Creating Learning Organizations: Growth Through Quality." It features Peter Senge, Michael Timpane, Robert Peterkin with "special guest" W. Edwards Deming.

The total program including local discussion is 2 1/2 hours and it costs $50.

Do you know anything about this program or its speakers (I've heard of Deming before, obviously)? Would you recommend attendance? I am writing a paper on applying PCT to school leadership and so this stuff (in addition to yours) could be relevant.--Gary

Date: Wed Feb 24, 1993 9:03 am PST
From: Dag Forssell / MCI ID: 474-2580
TO: Gary
Subject: Seminar

Peter Senge has written "The fifth discipline" the best popular introduction to systems thinking I have found. I have adapted some of his illustrations.

I would think you will get stimulated $50 worth by attending. Leadership and education go hand in hand.

Did I make sense in discussion of control systems a while back?

Enjoy!   Dag

Date: Wed Feb 24, 1993 10:11 am PST
Subject: Re: Language; designing robots vs. analyzing organisms

[Martin Taylor 930224 12:15] Bill Powers 930222.1300 ??

There's still a niggle. You agree with my comments on precise perception, but use them to come to the opposite conclusion than the one I draw. I'm mot sure where the difference lies, but let's wiggle a bit more.

>So I completely agree when you say that the limit of precision is
>seen as "precise." If we both keep this in mind we should be able
>to keep the domains of analogue modeling and statistical analysis
>properly separated.

I use the first statement to argue that the two domains are inextricably linked, not when it comes to design, perhaps, but when it comes to performance. It is in performance that the models are compared with living organisms and the plausibility of the models tested.

Perhaps the question is where we concentrate the interest. I agree with you that the statistical aspects are no great help in developing designs for "circuits" other than by eliminating implausible suggestions. For design, one can look at the gross requirements of a task -- just how would one control perception X, anyway? But then the statistical issues come to the fore. How fast could this circuit act? Could the robot lift the weight while not falling off the tightrope? Do living systems work faster or more precisely? or less so? Have we used (or omitted) an available information source that could account for the discrepancy.

To shift metaphors, I see the analogue design as providing the skeleton that determines the basic form, the statistics as providing the flesh and muscle. They can't be kept separate.

In the same posting, your reply to Bruce brings up once more a question that has been bugging me for some time: the "imposition" of structure on the hierarchy. What is the role of mimicry, in which one organism acts as it sees another do in similar circumstances. This seems as if the organism takes an S-R model of the other, and uses it to build perceptual functions and output links into the hierarchy. Only later is the function of this model discovered, and connected with reference levels that mean something. And in the same vein, teaching asserts a verbal structure that is somehow built into the hierarchy. Neither of these ways to pass socially important material seems to fit the standard description of reorganization. Both lack the essential element of randomness or of gradient search (which can be used for perceptual input functions).

A problem for me. How about you?

Martin

Date: Wed Feb 24, 1993 2:13 pm PST
Subject: Reading old posts and note a confusion

****************** FROM CHUCK TUCKER 930224 ********************

I realize that I am selecting out of context but there are some statements made by Martin and Bill that I think might make for some misunderstanding in others as to what PCT or HPCT is all about. I put forth as examples:

Powers 930217.1030 to Blom 930208

>You can't tell what a person is doing just by looking at what the person
>is doing. The test for the controlled variable helps you to understand
>what is actually being controlled (as opposed to what you logically assume
>is being controlled).

Taylor 930219. 19:00

>We can develop social conventions about actions (including language, as
>Bruce's lucid 4-parter shows) but we cannot develop social conventions
>about behaviour (in the PCT sense). Actions are are what we can see of
>what others are doing. We can't see what they themselves are "doing."

Now read those statements carefully and see if you understand what PCT or HPCT is about. Too most it would be quite confusing since they are trained in the objective positivistic perspective of human behavior. As we have noted it takes pages and pages of writing on the net to have most (we even "forget" it ourselves) grasp what is being proposed here. The model must take the perspective of the subject; the "role of the other"; the "empathetic view"; the view from the inside of the subject. It even allows for introspection. It does all of this while claiming to be more systematic and scientific that all of the extant theories of human behavior. Take the view of the reader of the above statements and just imagine how confused she must be when entering our conversation.

Finally, contemplate this statement by Powers 930217.1030

>The fact that all meanings are private -- together with the common assumption
>that they are objective -- explains most of the woes of the world.

Yes, reorganization is necessary for most to comprehend this model of human behavior.

Regards, Chuck

Date: Wed Feb 24, 1993 4:21 pm PST
Subject: Ashby vs. Powers

Hi,

I've been away from this list for a few weeks while I finished my preliminary literature survey for the upcoming contract with Martin Taylor. Basically, I've been grounding myself in some of the basic literature (although there's still tons of stuff I'd like to read). Now that I'm back, I've got literally hundreds of postings to catch up on, so forgive me if I'm responding to long dead discussions.

This posting is directed mostly towards Bill Powers, and his characterization of W.R. Ashby. As usual, Bill seems right on target until he starts talking about information theory, and he completely loses me (Ashby's formulation of control is mostly based on Shannon-style information theory).

Bill Powers (930118.1600) writes:

< Here's another gaggle of myths, this time from W. Ross Ashby, in
< _An Introduction to Cybernetics (New York: Wiley, 1966 (third
< printing, copyright 1963).

Its hard to believe you're talking about the same book! I didn't see any of the anti-error-control stuff you talk about.

[quoting Ashby]
<"_Regulation by error._ A well-known regulator that cannot react
<directly to the original disturbance D is the thermostat-
<controlled water bath, which is unable to say "I see someone
<coming with a cold flask that is to be immersed in me -- I must
<act now." On the contrary, the regulator gets no information
<about the disturbance until the temperature of the water (E)
<actually begins to drop. And the same limitation applies to the
<other possible disturbances, such as the approach of a patch of
<sunlight that will warm it, or the leaving open of a door that
<will bring a draught to cool it." (p. 222).

Bill, I would have thought you'd agree with this characterization of error control. Ashby is saying that error control cannot be perfect since it relies on detecting errors. If there is error, then obviously you are not controlling perfectly, by definition. This seems pretty straightforward and not at all in conflict with PCT. What's your beef?

<Note the implication that a compensating regulator might exist
<which, on seeing someone approach with a flask, could deduce that
<it contains cold water and is about to be immersed in the bath.

I did not notice any such implication. On re-reading these passages, I still do not. As far as I can tell, Ashby is not criticizing error control at all - it is the main thrust of his book in fact. The fact that he says it has "limitations" does not mean he thinks its a bad thing. Understanding the limitations of something is a step toward understanding its power. And just where is Ashby's big chapter on compensatory regulation? You imply this is his thrust, but I see mostly talk of error control in his book.

<Note also the unspoken assumption that merely from qualitative
<knowledge about a flask of cold water, a patch of sunlight, or a
<potential draught through an open door, the regulator could be
<prepared to act quanititatively... If, of course, such a
<thing were possible, the compensator would be much superior to
<any form of feedback controller. But such a thing is not remotely possible.

"Unspoken"? Are you sure he's making this assumption at all? I don't see it anywhere in the chapter on error control you are quoting from, and I do not remember any argument in Ashby's book for the kind of compensatory error control you are attributing to him.

<After doing through a series of diagrams, Ashby finally diagrams
<the true error-driven control system:

Your "finally" is probably justified. He goes through a rather long process, developing the notion of control from an information theoretic standpoint, in terms of open-loop communication channels. I would certainly bring in the closed loop earlier, but nonetheless, the closed-loop diagram is the whole point towards which his argument is leading. You make it sound like an afterthought. In fact, Ashby says of his closed loop error control diagram: "This form is of the greatest importance and widest applicability. The remainder of the book will be devoted to it." He adds of his earlier open loop diagrams (the "perfect" controllers): "(The other cases are essentially simpler and do not need so much consideration.)" This doesn't sound like the Ashby you describe.

<Now we get to a whole fountain of misinformation about control
<systems, a series of deductions that is just close enough to
<reality to be convincing, and just far enough from it to be utter nonsense.
<
<"A fundamental property of the error-controlled regulator is that
<_it cannot be perfect_ in the sense of S.11/3" (p.223)

Again, this is in line with PCT, no?

<[still quoting Ashby] ... So the more successful R is
<in keeping E constant, the more does R block the channel by which
<it is receiving its necessary information. Clearly, any success
<by R can at best be partial." (p. 223-224)
<
[Bill speaking again]
<This argument has apparently convinced many cyberneticists and
<others that the Law of Requisite Variety is more general than the
<principles of control, and in fact shows that control systems are
<poor second cousins to compensators when it comes to the ability
<to maintain essential variables constant against disturbance.

I think his argument is essentially valid and consistent with PCT. The controller gets its information in the form of the percept it is controlling. Thus, it can react only to error. The fact that it has detected error at all means that its control *cannot* be perfect. This is Ashby's basic point here: error-control is by nature imperfect. There is *nothing* here in conflict with PCT that I can see. Have any PCTers ever claimed that total perfect control was possible? By disagreeing with Ashby so heatedly, you are implying that you believe so. So tell me: how can a control system that controls against disturbances solely through reaction to error signals possibly control perfectly when the very presence of the error signal indicates imperfection? This is really the total extent of what Ashby is saying here.

<In fact this argument shows how utterly useless the Law of Requisite
<Variety is for reaching any correct conclusion about control systems.

The Law of Requisite Variety (or information) says simply that the complexity in the disturbance must be made informationally at least as low as the output channel capacity if the control system is going to successfully control against this disturbance. In terms of the above argument, R gets the information it uses to keep E constant through E itself, thus there is simply no way that it can keep E perfectly constant, since control involves the blockage of this channel. This is perfectly consistent with PCT.

<Having swept through this dizzying exercise in proving a falsehood, Ashby
<then grudgingly allows feedback control to creep humbly back into the picture:

Why do you say "grudgingly"? It doesn't sound at all grudging to me.

<Note also how the qualitative concept that error-regulated
<control must be imperfect is used to imply that it must be _more
<imperfect than compensatory regulation_.

Again, I'd really like to know where in the book he says this. Can you cite the chapters or sections? I just didn't get this from it at all.

<... He didn't know that the
<"imperfection" inherent in such systems can be reduced to levels
<of error far smaller than the error-reductions that any real
<compensating system could achieve -- smaller by orders of
<magnitude, in many cases, particularly cases involving human
<behavioral systems.

The compensatory regulators I see in Ashby are mathematical idealizations which he himself does not place much, if any, practical value on. They are a route for Ashby to understand the real issue: error control. I know you will probably say this is a crummy route, especially considering your low opinion of information theory, but it does not place Ashby in an opposing camp to PCT.

<Ashby's entire line of reasoning about feedback control in _An
<introduction to cybernetics_ is spurious. Yet Ashby has been
<revered in cybernetics and associated fields for 40 years as a
<deep thinker and a pioneer. His Law of Requisite Variety has
<nothing at all useful to say about control systems -- and in fact
<led Ashby to a completely false conclusion about them -- yet it
<is still cited as a piece of fundamental thinking. Whether Ashby
<originated these misconceptions or simply picked them up from
<others I don't know. One thing is certain: he did not get them
<from an understanding of the principles of control.

I think you've set up a straw man, but I'm willing to consider that I just misunderstood Ashby, seeing him through a PCT-lense. Again, it would help to have some specific references to places in the book where he makes the case for compensatory control.

===
Allan Randall, randall@dciem.dciem.dnd.ca
NTT Systems, Inc. Toronto, ON

Date: Wed Feb 24, 1993 9:53 pm PST
Subject: Statistics; Ashby

[From Bill Powers (930224.1730)] Martin Taylor (930224.1215) --

>How fast could this circuit act? Could the robot
>lift the weight while not falling off the tightrope? Do living
>systems work faster or more precisely?

The limitations placed on control by physical properties of the system and its environment are seldom, in robotics or servo design, statistical in nature. The bulk of servomechanism texts is devoted to the subject of stability -- that's what all those Laplace transforms and Bode diagrams are about. The underlying problems are in the field of differential equations, and assume complete lack of noise or randomness in the signals. It isn't that engineers fail to take random effects into account, but that the normal case treated first is the one using noise-free variables. If the system isn't stable using smooth variables, it won't be stable when noise is present, either. And conversely, if it's stable in the absence of noise, it will be stable in the presence of noise, too, although it may not control as well.

The boundary between statistical treatments of signals and continuous-variable treatments is found where signals or signal variations of interest get so small that their discrete-impulse nature begins to make a difference. The normal range of the signal emitted by a single neuron is (I would estimate) approximately between 1 and 1000 impulses per second. An ordinary perception probably entails pulse rates of several hundred per second. As most perceptual channels probably involve redundant parallel paths, the average rate is probably higher than this.

If we assume that nerve signals are random (Poisson-distributed) variables, then the signal-to-noise ratio goes about as the square root of the pulse rate for short sampling times. This would lead to a maximum SNR of about 30 for a signal running at 1000 pps, and 10 for a signal at 100 pps. However, the dynamic range is probably greater than the amount implied, because signals are not generated at random by neurons but as fairly regularly space impulses. The real noise level is in the deviation of the impulse trains from perfectly uniform repetition rates. I don't know the actual numbers here, but the variations are clearly less than the Poisson distribution would suggest. Judging from JND experiments, for some sensory channels the noise in the signal may be around 5 +/- 2 percent of the signal magnitude.

This is more or less what we see in tracking experiments. The model, which contains no noise, reproduces the wobbles of the subject's handle within 3 to 5 percent, reproducing not just the same statistical magnitude of wobbles, but tracing the very same wobbles; this shows that the wobbles are not statistical.

When the model handle behavior is subtracted from the subject's handle behavior, we are left with a difference-trace. Most of the remaining wobbles are relatively random-looking and of a much higher frequency than those that the model reproduces. I would estimate that the truly random differences between the real and model behaviors amount to 2% or less of the peak-to-peak handle excursions (there are no doubt systematic wobbles that the model doesn't reproduce). I would quess that this noise level is showing the underlying statistical noise in the system.

All this is by way of preparing to comment on this:

>To shift metaphors, I see the analogue design as providing the
>skeleton that determines the basic form, the statistics as
>providing the flesh and muscle. They can't be kept separate.

If a statistical analysis shows that the signal variations have a random component of less than 5 percent, this by no means guarantees that a stable control system exists, or that it will control well. Even reducing the noise level and uncertainty to zero will not cure an unstable system.

As I see it, the statistical treatment becomes important when the basic uncertainties in the system begin to be a significant fraction of the magnitude of a reference signal -- say 5 or 10 percent. Below that level, the noisiness in the error signal probably has little impact on behavior, especially considering that higher-level systems can combat any low-frequency errors that leak through by adjusting the lower-level reference signals.

My impression is that you consider the uncertainties in perception, comparison, and action to be much larger than I do. I think that a large part of our pulling and tugging over this subject has to do with the assumed magnitude of uncertainties. It could also result from your interest in higher-level systems about which I can say little in terms of experimentation.

>In the same posting, your reply to Bruce brings up once more a
>question that has been bugging me for some time: the
>"imposition" of structure on the hierarchy. What is the role
>of mimicry, in which one organism acts as it sees another do in
>similar circumstances.

I use the term "imposition" in connection with the construction of higher-order variables as functions of lower-order ones. Without a higher level of perception, a set of lower-level perceptions simply coexist. Without any added considerations, one could in principle construct a new perception that is any conceivable function of the existing ones. The behavior of the new higher-level control system would then force the lower-level world into a condition that makes this arbitrary function of the lower-level variables assume an arbitrary desired value. Thus the higher-level system imposes a degree of order on the lower-level world.

Mimicry is a delicate question, because the perceptions that one organism has of another's actions are obtained from the wrong point of view. If I say to you, take off your glasses and hold them in the same orientation I'm holding my glasses, you're seeing my glasses from the side opposite to the side on which I'm seeing them. Yet you orient the glasses relative to yourself so you see the same side of them that I see of mine (and the side of mind that I don't see). You don't simply hold your glasses so they give the same perceptual image you're getting from mine.

It's much easier to explain mimicry when it involves repeating some effect on the environment that's more or less independent of point of view. If I put my fork "in" my glass of water, you can put your fork in your glass of water by recreating the "in" relationship between objects of the same appearance. This is much more like simply repeating a perceptual experience that you remember. Other kinds of mimicry, including paraphrasing, involve some complex transformations and imagination.

If one organism builds an S-R model of the other, as you suggest, this would amount to repeating a relationship between the antecent event and the subsequent action. I could imagine reproducing the "causality" relationship without actually reproducing the specific act from the viewpoint of the other.

>Neither of these ways to pass socially important material seems
>to fit the standard description of reorganization. Both lack
>the essential element of randomness or of gradient search
>(which can be used for perceptual input functions).

I agree; the processes involved are much too systematic to be called reorganization. If I had to try to model mimicry, I'm sure that at some point I'd have to depend on program-level processes for deducing the appearance of the world from a point of view other than my own. Reorganization might be needed to come up with the necessary transformations, but once they have been learned they are simply the execution of a program.

Chuck Tucker (930224) --

>... there are some statements made by Martin and Bill that I
>think might make for some misunderstanding in others as to what
>PCT or HPCT is all about.

Well, it's fun sometimes to speak in riddles just to see if someone else can figure out the solution. You can't tell what a person is doing (mailing a letter on a rainy day) just by watching what the person is doing (putting on a raincoat).

Martin's comment also is a riddle, to the effect that we see other people's outputs, but can't see the perceptions they're controlling. What we mean by "doing" is the other's actions; what the other person means by the same word is the perceptual situation brought about by those actions. You see me wiggling a steering wheel and causing the car to swerve; I see myself making the front wheels miss a pothole.

As long as the reader gives us the benefit of the doubt by assuming that somehow such paradoxical statements make sense, the riddles serve a heuristic purpose. They challenge the reader to find the sense in them. Of course if the reader takes everything literally, we just seem to be saying one thing and then the opposite.

You might say that the riddles are a screening device.

Allen Randall (930224) --

RE: Ashby.

Me quoting Ashby --

<"_Regulation by error._ A well-known regulator that cannot react
<directly to the original disturbance D is the thermostat-
<controlled water bath, which is unable to say "I see someone
<coming with a cold flask that is to be immersed in me -- I must act now."

You:

>Bill, I would have thought you'd agree with this

>characterization of error control. Ashby is saying that error
>control cannot be perfect since it relies on detecting errors.
>If there is error, then obviously you are not controlling
>perfectly, by definition. This seems pretty straightforward and
>not at all in conflict with PCT. What's your beef?

What I was objecting to was that Ashby was contrasting the operation of a compensatory system with the way an error-based control system works by saying the control system could not act on the basis of a threat, and using an example of a sort in which NO compensatory system could behave successfully. He seemed to be claiming that a compensatory system could prevent an error in bath temperature by seeing that someone is coming with a cold flask, and acting NOW. In fact, that would be more likely to cause error than to prevent it, unless the system could somehow make quantitative guesses as to how much heat the flask and its contents were going to absorb from the bath, and could calculate just how many BTUs to add to the bath, and EXACTLY WHEN, to prevent any error. While Ashby is being critical about the abilities of control systems, he fails to be equally critical about what his proposed system could possibly accomplish under the stated conditions. This is how straw-man arguments get started.

Me:
>>Note the implication that a compensating regulator might exist
>>which, on seeing someone approach with a flask, could deduce
>>that it contains cold water and is about to be immersed in the bath.

You:
>I did not notice any such implication.

Ashby said
<.. the thermostat-controlled water bath, which is unable to say "I see someone
<coming with a cold flask that is to be immersed in me -- I must act now."

I assumed this to imply that there IS a kind of system that could achieve control by anticipating a disturbance of this sort. As this is the kind of system he had been (prior to this point in the argument) presenting, I took this to be an implication that a disturbance-compensating system COULD do what the error-based system could not do.

>And just where is Ashby's big chapter on compensatory regulation? You
>imply this is his thrust, but I see mostly talk of error control in his book.

The basis runs from p. 140 to p. 160 (_An introduction to cybernetics_) in the chapter called "transmission of variety." The application to the subject of control begins in Chapter 10, where the Law of Requisite variety is developed, and arrives at the compensating system (as I call it) on page 210 in chapter 11.

```
     ->T
   /   ^   \
 D     |     -->E
   \   |
    -> R
```

Ashby pictures D as a disturbance, which acts on an intermediate device T which in turn affects an essential variable E. A regulator R is provided, which senses the state of D, and as a result acts on T to compensate for the effect of D on T, thus leaving the essential variable E undisturbed.

>I do not remember any argument in Ashby's book for the kind of
>compensatory error control you are attributing to him.

Perhaps the above diagram will refresh your memory. The effect of the disturbance D on E through the path D-T-E is compensated for by the effect of the disturbance on E by the path D-R-T-E.

For this kind of arrangement to be capable of "regulating" E, the effects of D via the two paths must quantitatively cancel (without any feedback to indicate when cancellation has happened). Any nonlinearity in the path D-T must be exactly compensated by a corresponding nonlinearity in the path D-R-T. The net effect on the essential variable E by the path T-E must be small in comparison with the effect D-T and the compensating effect D-R-T (if regulation is truly achieved). Thus the regulation takes place through a small difference between two much larger effects on T. This places severe precision requirements on the regulator R and the paths D-R and R-T. It also requires that there be no undetectable sources of disturbance acting directly on E or T. In short, this scheme has all of the defects of compensating systems that I have laid out previously. In practice it is not likely to achieve anything like the degree of error-correction that could be achieved by a system that sensed E directly and compared the sensed state against the desired state.

Ashby definitely is assuming that a regulating scheme like the one above could actually achieve significant regulation. It could not, in the real world.

On page 213, you will find an elaboration on the above scheme at the beginning of the section called "control". The elaboration consists of a box C that acts on R to establish a "target" for the state of E. The diagram is

```
     ->T
   /   ^   \
 D     |     -->E
   \   |
    -> R
      /
 C -->
```

Now the regulator must contain different transformations that will act through T to produce different states of E. The input from C selects which of these different states will be achieved. Ashby says "Thus the fact that R is a perfect regulator gives C perfect control over the output, in spite of the entrance of disturbing effects from D." (p. 213-214).

Ashby's fatal assumption was that R could in fact be anything close to a perfect regulator. In comparing the behavior of this kind of system against a closed-loop system, he compared a compensating system of highly exaggerated precision against an error-based control system with grossly-underestimated capabilities.

>>"A fundamental property of the error-controlled regulator is
>>that _it cannot be perfect_ in the sense of S.11/3" (p.223)

>Again, this is in line with PCT, no?

The problem is that his is a qualitative statement: either regulation is perfect or it isn't. If he had considered real physical situations in his diagrams above, he would have had to make the same statement about the compensation approach: it cannot be perfect.

If, on the other hand, he had considered both kinds of system quantitatively, the judgment would have been that while control by error and control by compensation are both necessarily imperfect, control by compensation is inherently far more imperfect.

Consider Ashby's second diagram with a small addition: a random disturbance of magnitude d added to the disturbance box output to T. This would be equivalent to the regulator R having slightly imperfect knowledge of the state of D. The basic functions are, as Ashby assumed, perfect otherwise.

The effect of d will be passed directly to E, because the regulator R does not detect this extra little disturbance.

Now assume an error-based control system of the same type:

```
             d
              \
             D --> T --> E
                   ^     |
                   |     v
           C -->   R <--
```

We assume that the output of R is G* (C - E), where G is the amplification factor. Note that R does not detect the state of D or d at all. The result will be that E = C* G/(1+G). If G is 1000, E = 0.999C, with an error of 0.001*C. The effect of the total distubance d+D on E is (d+D)/1001. If the change of E by 1/1000 of the total disturbance is considered too imperfect, the gain can be raised to a million, or a billion, with corresponding decreases in the effect of the total disturbance. Yes, the control system is imperfect, because some error is required to produce output. However, the amount of the imperfection is so small that in most circumstances we would not quibble about it.

>I think his argument is essentially valid and consistent with
>PCT. The controller gets its information in the form of the
>percept it is controlling. Thus, it can react only to error.
>The fact that it has detected error at all means that its
>control *cannot* be perfect.

No, it cannot be perfect. However, it can be so close to perfect that no human eye could tell the difference. This is what Ashby didn't know, because he had no experience with good tight control systems. To him, the necessity for an error signal made him throw up his hands and say, "oh, well, then...". If he had understood how small the error signal can be in a good control system, he would never have talked about "imperfection" in such loose terms. The imperfections in real control systems can be orders of magnitude less than those in compensatory systems.

It isn't that I'm arguing against Ashby's concept of closed-loop control. I'm just saying that he never went into the subject deeply enough to see how such systems really work, and therefore didn't see that a closed-loop regulator is ALWAYS better at regulation than any open-loop regulator of the types he diagramed. He was overimpressed by the "imperfection."

>The Law of Requisite Variety (or information) says simply that
>the complexity in the disturbance must be made informationally
>at least as low as the output channel capacity if the control
>system is going to successfully control against this disturbance.

But this is just a complicated and abstract way of saying that the output of the control system must quantitatively oppose the effects of the disturbance. The condition you state is more general than the actual requirements, and furthermore doesn't help in designing a system that could actually produce the described result. "Complexity" isn't the problem; error is. This is one of those true but useless statements.

My biggest guffaw at Ashby's reasoning came on p. 223-224, where he goes through the logical error that practically every neophyte in control theory goes through. "Control systems correct error. But their actions are driven by error, so if there's no error, they can't act, and if they can't act, they can't correct error." This is what you get for reasoning your way sequentially and qualitatively around the control loop instead of solving the simultaneous equations. The guffaw came when Ashby, having made this elementary conceptual error, proceeded to show how it follows from the Law of Requisite Variety.

>>Note also how the qualitative concept that error-regulated
>>control must be imperfect is used to imply that it must be
>>_more imperfect than compensatory regulation_.

>Again, I'd really like to know where in the book he says this.
>Can you cite the chapters or sections? I just didn't get this
>from it at all.

I've already noted the sections, but here they are put together:

P. 213: "Suppose now that R is a perfect regulator." This implies that R can be perfect, at least in principle. The argument that follows is based on a perfect regulator.

P. 223: "A fundamental property of the error-controlled regulator is that _it cannot be perfect_." So an error-based control system can't be perfect, even in principle. The arguments that follow are based on the idea that error-controlled system are inherently -- and significantly -- imperfect.

The only way for R to be a perfect regulator is for the variables it handles to be logical or binary and the calculations involving them error-free. The examples Ashby uses always use small whole numbers or logical propositions, in which small errors never occur. A disturbance of 3 units is canceled by a regulator output of -3 units: perfect regulation. If we considered real regulators, even those involving binary variables, we would find nonzero errors, and the in-principle perfection would vanish.

The imperfection of the control system that is seen depends on the time scale used (but that is true of compensating regulators, too). It also depends on what you imagine to be the limits on loop gain. I can't find the reference right now, but at one point Ashby "proved" that the maximum loop gain of a stable negative feedback system (implemented in discrete operations) is less than -1. Perhaps that is why he thought errors had to be large enough to be bothersome. In principle, a stable control system can actually have a loop gain of negative infinity, at which point the error drops to zero, and control is perfect. In practice, loop gain must be traded off against error-correction speed, so control systems actually can't be perfect either.

If Ashby had made his comparisons in similar terms, he would not have made these mistakes.

>The compensatory regulators I see in Ashby are mathematical
>idealizations which he himself does not place much, if any,
>practical value on. They are a route for Ashby to understand
>the real issue: error control.

But when Ashby finally got to error control, where did he go with it? To "The Markovian Machine," another exercise in manipulating discrete variables and small integers -- and by necessity, low-gain control systems. The best I can say about these machines is that they may have some relationship to reorganization. The loop gains in his diagram are always kept carefully less than 1.

On p. 236 he finally gets to control systems with continuous variation, mentioning temperature control, pH control (homeostasis), and power amplification, for a total of about 4 pages of text. Having covered what he knew about that subject, he then goes on to games and strategies -- back to discrete variables again. Then it's off to regulation of very large systems (discrete variables again), and finally an incoherent chapter on "amplifying regulation."

Ashby was actually an idol of mine after I read "Design for a brain" shortly after it was published. He said many true and useful things about control and negative feedback in that book, at least in the early parts, and gave me many useful ideas. It probably took me 20 years to realize how little Ashby actually knew of a technical nature about control. Mostly he was following his nose and writing down what seemed intuitively reasonable to him, and picking up on whatever ideas were trendy at the time. He was, after all, a psychiatrist, not an engineer. He was probably a pretty good mathematician. From all accounts he was a terribly nice guy.

Best to all, Bill P.

Date: Thu Feb 25, 1993 9:53 am PST
Subject: Re: Statistics; Ashby

[Martin Taylor 930225 12:00]
(Bill Powers 930224.1730)

I still don't know whether our disagreements on statistics within PCT represent a gulf or an insignificant rill. There's nothing that I want to complain about in your discussions of stability factors and system design. But I don't think that the size of the statistical uncertainty is where we differ.

>It isn't that engineers fail to take random effects into account, but that
>the normal case treated first is the one using noise-free
>variables. If the system isn't stable using smooth variables, it
>won't be stable when noise is present, either.

Fine. I think we have no problems in agreeing along these lines.

The point that looms ever larger to me, and that you diminish, is that any control system will produce output to the degree that it can sense error. And that depends on the statistics. Whether it is of any value to continue pressing this point, I don't know. It is at the heart of the information -> PCT paper, so maybe it is best to leave off the discussion until I can get around to continuing that.

Allan gave me a xerox or one or more chapters of Ashby, which I haven't yet read. The following depends entirely on Bill P's characterization of Ashby, not on Allan's.

```
>The application to the subject of control begins in Chapter 10,
>where the Law of Requisite variety is developed, and arrives at
>the compensating system (as I call it) on page 210 in chapter 11.
>
>                       ->T
>                    /    ^   \
>                  D      |     -->E
>                    \    |
>                      -> R
>
>Ashby pictures D as a disturbance, which acts on an intermediate
>device T which in turn affects an essential variable E. A
>regulator R is provided, which senses the state of D, and as a
>result acts on T to compensate for the effect of D on T, thus
>leaving the essential variable E undisturbed.
```

As I look at Bill's description of the compensating system, it seems very like the imagination loop. "Controlling" by imagination in the absence of sensory input is not going to be very effective, as Bill many times has pointed out. But if the model of the world incorporated in R is even moderately good, the error will usually be smaller than if R is absent, and its evaluation will require less sensory information. The control bandwidth can be wider or control can be more accurate for the same bandwidth.

I'm not commenting on Bill's characterization of Ashby's understanding, at least not until I get around to reading Ashby (if I do). But we are currently trying to understand how world models and imagination are incorporated into controllers, and it seems as though Ashby's circuits fit well into our current discussions.

>>The fact that it has detected error at all means that its
>>control *cannot* be perfect.
>
>No, it cannot be perfect. However, it can be so close to perfect
>that no human eye could tell the difference.

In fact, it can be that perfect AND NO MORE SO, if the control is of a visual perception. (A question of statistics).

Martin

Date: Thu Feb 25, 1993 10:49 am PST
Subject: EEG mind reading, statistics

[From Rick Marken (930225.0900)] Martin Taylor (930223 0:30) --

>>I think the description of the article shows the naivete of
>>people's concept of behavior; the article is clearly based on
>>the proposition that the brain (EEG waves) commmands what we
>>do (type letters). All the snazzy physiological equipment in the
>>world cannot overcome such a fundemental error.

>Well, many years ago, I did see a demonstration at McGill University in
>which the subject "thought" a square on a screen (a cursor, if you like)
>into one of the four corners of the screen, using the EEG. It worked,
>and to say on purely theoretical grounds that it didn't is a rather
>fundamental error. I know nothing of more recent work in the area, but
>it doesn't seem outrageous to me that people could "think" words onto
>?a screen in the same way, if the vocabulary were small enough.

I guess I wasn't clear. I have no doubt that EEG waves can be transformed into letters. I would not even be surprised if a person could learn to use their own EEG outputs to produce (relatively reliably) particular letters on the screen. In this case, the EEG would be part of a control loop in what would be called a "biofeedback" situation. I imagine that is what was going on at McGill; the EEG output was hooked up in some way to an oscilliscope so that a certain kind of variation in the EEG would move the square horizontally and another would move it vertically. I have done biofeedback control (as a subject -- courtesy of David Goldstein) and I did a pretty good job of controlling the size of a square on the screen using my GSR. It wasn't real tight control (unfortunately, I never controlled against disturbance so I have no quatitative idea of how good the control actually) but it was control nevertheless.

So I was not rejecting a finding on theoretical grounds; I was rejecting what I read as the interpretation of the finding -- that certain components of the EEG signal actually corresponded to "letters" in some sense -- as perceptions or "command" outputs.

By the way, I don't think I would EVER rule out an OBSERVATION on theoretical grounds; if I saw someone lift a table off the ground by waving their hand over it, I would accept that observation; it's the interpretation of the observation that I might question (that the table went up simply because the hand was waved above it). This is all I meant about the EEG result; I believe it; I just don't believe the (implied) interpretation (which I was gleening from the very brief description) that some quantitative measure of EEG is the neural correlate of letter perceptions or letter command signals.

Martin Taylor (930224.1215) --

>To shift metaphors, I see the analogue design as providing the
>skeleton that determines the basic form, the statistics as
>providing the flesh and muscle. They can't be kept separate.

Then why are we able to build models that make point by point predictions of long stretches of behavior with less than 2% error -- while never considering statistics AT ALL? Could you give a concrete example of how the incorporation of statistical considerations into the PCT model could improve it's performance. Much of this discussion of statistics seems somewhat religious to me -- statistics just MUST be important. Well (as Eliza Dolittle said) SHOW ME! Why do we need statistics?

This is an important topic to me because, in my eary days as a PCT rebel (when I thought facts and precise models would be enough to convince people that PCT was a better idea) I ran into the most remarkable phenomenon; people who dismissed the results of my experiments and modelling because the there were NO STATISTICS. These people (colleagues as well as journal editors and reviewers) believed that inferential statistics are just a necessary part of psychological science: statistics is to psychological science as eucharist is to catholic mass. No statistics, no science. The fact that the correlation between model and subject data was on the order of .997 (and, thus, required no inferential test to determine the probability that there might not "really" be a relationship between subject and model data) was taken as evidence that the results were trivial (I guess there was no more variance to account for so there was obviously nothing else to do).

I believe it is this religious treatment of statistics in psychology that makes it impossible for conventional psychologists to see any value in PCT. I mean, if you are presenting data that is just about as good as what you get in a physics lab and the "lab teacher" is saying "that's not very interesting data because it's not the kind we like -- which is noisy and difficult to interpret without statistics -- then what can you do? Do we have to get junky data in order to prove that we are a real science? Jeeez.

Chuck Tucker (930224) --

>... there are some statements made by Martin and Bill that I
>think might make for some misunderstanding in others as to what
>PCT or HPCT is all about.

I like Bill Powers (930224.1730) comment on this one:

>You might say that the riddles are a screening device.

By the way, that "Mind and Healing" show, though fairly silly some- times, is pretty good and it sure does speak to issues of intrinsic interest to PCT. In fact, I just happened to be reading the Emotion chapter in LCS II before I saw the show; I think that this is an excellent chapter -- and one that a lay audience can resonate to. Like all aspects of our experience, the inference is that illness perceptions depend on a "body". I think Moyers' show would be more interesting if it were based on the question "what does the body have to do with the mental experience that is illness?" instead of "what does the mind have to do with bodily illness?". Maybe Eastern psychology and life science would be more receptive to PCT?

Best Rick

Date: Thu Feb 25, 1993 12:29 pm PST
Subject: Slowing vs. Perceptual Lag

[from Gary Cziko 930225.1700 GMT] Bill Powers (930224.1730) said:

>In principle, a stable control system can
>actually have a loop gain of negative infinity, at which point
>the error drops to zero, and control is perfect. In practice,
>loop gain must be traded off against error-correction speed, so
>control systems actually can't be perfect either.

It's taken me a while, but thanks to your tutoring on the net and your Demos 1 and 2 I am developing what I think is a pretty good understanding of (at least) simple proportional control systems.

But your comment reminded me of a problem I still have. I can understand the relationship between loop gain and slowing. I played with the demos and seen what happens if the loop gain is too high for the slowing (or integration) factor. But what I don't understand is the difference effects of the slowing factor (which is actually the how quickly error is turned into action) and perceptual lag (which is how quickly changes in the physical, controlled variable are reflected in changes in the perceptual signal.

Why is it that slowing from error to action is "good" (at least in the sense that it provides stability with high loop gain) while perceptual lag is not good. Intuitively, perceptual lag is bad because if I don't see the result of my actions until later I will continue overcorrect. And yet, both output slowing and perceptual lag just seem to add up to an environmental delay in the closed control loop.

Can you understand my dilemma? What am I missing here?--Gary

Date: Thu Feb 25, 1993 1:26 pm PST
Subject: Re: EEG mind reading, statistics

[Martin Taylor 930225 15:30] Rick Marken 930225.0900

OK on the EEG. It is getting scary how often we seem to be agreeing these days.

But there's compensation on statistics.

As I read your posting, your objections to the statistical analysis they teach in psychology grad school have carried over into a generalized objection to thinking statistically. You start your posting by talking against my uses of statistics within the circuits of a hierarchic control system, and without missing a beat you go into objections about the use of statistics by an analyst.

>Then why are we able to build models that make point by point
>predictions of long stretches of behavior with less than 2%
>error -- while never considering statistics AT ALL?

Why shouldn't you? If your circuits are correct, controlling the appropriate perceptions, with the same kinds of limitations as the human, 2% might be considered a horrendous error. Correctness of conception more or less guarantees that you will be pretty close. If you limit the task difficulty, the statistical limitations won't show up. But if you make the cursor in a tracking task move too fast or with too high a bandwidth, the human will lose control. A model with infinite bandwidth won't. You have to limit the model in the same way that the human is limited, if you want to maintain your prediction accuracy under such conditions. That limit is in the speed and resolution of perception, as you yourself pointed out in one of your papers (Hierarchy of perception?). That limit is based on statistical considerations.

I don't care whether the models fit the data to an accuracy of one part in a million or a billion. The argument is that the performance of the model AND of the human is limited by the information available to the perceiving system. I don't see why this presents you or Bill P with a problem. It's almost a restatement of the fundamental position of PCT: you can't control what you can't perceive.

> I ran into the most remarkable phenomenon; people who dismissed the results
>of my experiments and modelling because the there were NO
>STATISTICS. These people (colleagues as well as journal editors
>and reviewers) believed that inferential statistics are just a
>necessary part of psychological science: statistics is to psychological
>science as eucharist is to catholic mass. No statistics, no science.

I totally sympathize with that side of your posting. And I have said so on other previous occasions. I wish you wouldn't keep trying to tar me with that brush.

I, too, have had papers returned for revision on these grounds, and one finally rejected when the editors wouldn't accept my arguments as to why the proposed statistics were useless garbage. In all these cases, I refused to include statistical analyses I thought were inappropriate, and that includes all significance testing (in my younger days, I did capitulate sometimes, though). On the other hand, I have published papers in which the statistics were central to the process under consideration--the person was, for example, affected by the variability of something rather than by its value. To look at the world with a stiatistical eye is simply to ask "what can be perceived."

I use statistics IN THE MODEL, not in the analysis (unless the situation warrants, which it sometimes does--we might have a legitimate disagreement on that last phrase). I have spent my academic life since half-way through grad school trying to persuade people of the evils of significance testing. So, if you don't mind a little repetition of past pointless rebuttals, I do feel a little insulted by:

>I believe it is *this* [my emphasis, MMT] religious treatment of statistics
>in psychology that makes it impossible for conventional psychologists to see
>any value in PCT. I mean, if you are presenting data that is just about as good

>as what you get in a physics lab and the "lab teacher" is saying "that's not
>very interesting data because it's not the kind we like -- which is noisy
>and difficult to interpret without statistics -- then what can you do? Do we
>have to get junky data in order to prove that we are a real science? Jeeez.

As well as not believing that this is part of the reason for PCT being ignored, I don't like it being applied to me.

Martin

Date: Thu Feb 25, 1993 1:50 pm PST
Subject: correction (Ashby)

[Martin Taylor 930225 16:30] Martin Taylor 930225 12:00

I said:
>As I look at Bill's description of the compensating system, it seems very
>like the imagination loop.

On re-reading, I realize this is subject to misinterpretation. What I mean is that the function R looks like, and has much the function of, the model that we seem to be developing as a necessary component of the imagination loop. The quoted sentence couldn't be right as it stands. The revision might also be wrong, but it's closer to what I should have said.

Sorry about that. Martin

Date: Thu Feb 25, 1993 2:42 pm PST
Subject: Re: statistics

[From Rick Marken (930225.1400)] Martin Taylor (930225 15:30) --

First, I want to apologize for implicitly "taring you with the brush" of affiliation with the "inferential statistics" religion. I forgot that we went around on that before; you are definitely a hero in the fight against requiring the shibbolith of statistical significance tests as a test of scientific significance in psychology.

Second, with respect to statistics and control you say:

>I use statistics IN THE MODEL, not in the analysis

Let me suggest that if there is any disagreement between us on this matter, it is one of emphasis, probably not of substance. I think Bill P. and I agree with you that, at the "micro level", say, at the level of neural firing patterns, stochastic process are involved. And I don't doubt that the control model could be improved somewhat by trying to take these stochastic processes correctly into account. But, at the level at which we typically look at control phenomena, stochastic phenomena just don't seem to play a big role; they probably contribute VERY LITTLE to the unexplained 1-2% of the variance. So the search for stochastic forms that might improve the predictions of the control model looks like a lot of effort for very little payoff.

But I think that you are making a case for the importance of stochastic processes "in principle"; that stochastic process place limits on control or require control. That's OK with me; I guess I just don't see ANY principle as being particularly important unless I can see how it is to be translated directly into a rule of model operation. That's why I would like to know 1) HOW you build stochastic processes into the control model and 2) what it buys you (in terms of increased prediction).

Best Rick

Date: Thu Feb 25, 1993 3:30 pm PST
Subject: Error statistics; thinking letters; feedback lags

[From Bill Powers (930225.1400)] Martin Taylor (930225.1200) --

>The point that looms ever larger to me, and that you diminish, is that any
>control system will produce output to the degree that it can sense error.

As you have said, part of communication is trying to make a model of the other person. I'm trying to understand -- model -- the mental model behind what you said above. The impression I get is that you consider noise to be a rather large component of the error signal, especially at zero error.

The problem here may be an unconscious change of scale when considering behavior under low-disturbance conditions or when the controlled variable is being maintained at a value near zero. Suppose we consider an arm movement in the horizontal direction about the shoulder, where there is a range of about 180 degrees of movement. This defines the range over which a position reference signal can be varied while control continues.

If the full range of magnitude of the reference (and perceptual) signal corresponds to 180 degrees or say 80 cm along an arc, how big is the error signal IN COMPARISON WITH THE RANGE?

When I hold an arm out straight and try to keep the forefinger still, the finger wobbles laterally around by no more than 2 mm (for a while anyway). This is 1/400 of the total range. The question is, how are we to speak of that 2mm of noise in the control system? If we compare the average error, which is less than 2 mm, with the RMS wobbles, then it will seem that the error signal consists MOSTLY of noise, with very little signal. My impression is that this is how you have been thinking about it.

On the other hand, if you compare the position wobbles with the total range of position, you find that the error signal fluctuations are only about 0.25 percent of the range of the perceptual and reference signals. So clearly, even if the error signal were 100% noise, the noise would not cause any appreciable error on the scale of the whole range possible. The magnitude of the error fluctuations due to noise is small on that scale. But if you magnify them until you can see the remaining average position error, they look enormous.

This isn't quite the end of it. In the example, when the arm stops moving there is essentially zero output force, reflecting the fact that there are no lateral disturbances. The position error signal averages essentially zero.

Now, however, if we apply a lateral force to the arm while the person continues to maintain the same finger position, the arm forces will become strongly nonzero. We can deduce that the error signal has also departed from zero, sufficiently to oppose the disturbing force. By observing the finger movements, we can now estimate the relative magnitude of the error signal and the underlying noise level. The finger movements will now be somewhat larger, probably, but not much.

If the random fluctions in error signal were still significant in comparison with the average magnitude of the error signal, we would expect to see significant fluctuations in the force opposing the disturbance. But we don't. The force is essentially steady. The finger wobbles are not importantly different from what they were in the absence of the disturbance. The systematic component of the error signal is now obviously much greater than the noise component.

I think you may have been doing the modeling equivalent of dividing by zero. By selecting conditions under which the average error is zero, you artificially inflate the importance of random noise in the error signal. Under other conditions, the error can depart from zero by many times the magnitude of the noise fluctuations (still remaining reasonable small in terms of the whole range of control). Suddenly the signal-to-noise ratio of the error signal looks high instead of essentially zero.

None of this vitiates the underlying statistical analysis of noise, and how the integration times in the control system reduce its effects. I suppose that there's a way to translate that into terms of information theory, too. The real question here is how you visualize the relative sizes of these effects; the relation of the noise level to the normal range of systematic error signals, and those to the total range of control of the perceptual and reference signals. If you picture the noise signals as being comparable to the size of normal error signals, noise will seem to dominate behavior. But if the noise signals are only a small part of the range of error signals, then they are not so important.

>As I look at Bill's description of the compensating system, it
>seems very like the imagination loop.

Did you look carefully at Ashby's diagram? There are no closed loops in it.

>In fact, [control] can be that perfect [to the acuity limit] AND NO MORE SO,
>if the control is of a visual perception. (A question of statistics).

I was being a bit more general: Mario Andretti could keep a car on a path so straight that I could see no errors, although he obviously could.

Rick Marken (930225.0900) --

My impression was that the EEG waves were run through a pattern generator that would type out the letter, with the person controlling the letter typed just by imagining it. Would someone who has actually read the article or seen the demonstration tell us what this is really about?

Pretty useful "biofeedback" for a lot of people, if it really works.

Gary Cziko (930225.1700 GMT) --

>But what I don't understand is the different effects of the
>slowing factor (which is actually the how quickly error is
>turned into action) and perceptual lag (which is how quickly
>changes in the physical, controlled variable are reflected in
>changes in the perceptual signal.

I presume you're talking about the experiments at the end of Demo 2.

In a system that has a slowing factor in the output function, gain and speed are traded off just as you understand them to be. In the demo program there is an added factor, a transport lag in the input function, that is introduced IN ADDITION to the output slowing factor.

The reason is that when a lag AND a slowing factor are in the loop, the system gets a little less stable -- it starts to overshoot after a sharp disturbance. Control gets a little WORSE. When, in the second part of the experiment, you adjust a perceptual lag, you're making the model's performance worse, so it tends to overshoot a bit. This makes it fit the person's performance BETTER.

That's the real reason for putting the lag in (aside from the fact that perception probably does actually involve some lags). The model with only a slowing factor in its output behaves too smoothly; it controls too well. As you increase the lag, you can see that the model starts to show more rapid movements. These are actually making its control worse, but they are making the difference between the model and the person smaller. So this tells us that we're at least in the right ballpark; the person, too, must contain something like the added lag.

There's a technical reason for two integrations, or one integration and one lag, being worse than one integration. With one integration, the largest phase shift that can occur with a sine-wave disturbance of high frequency is 90 degrees. A 90- degree lag in the feedback can reduce the effectiveness of the feedback, but it can't turn it into positive feedback.

When you have two lags, however, the first lag, for high frequency disturbances, can approach 90 degrees of feedback lag, and the second lag can add another 90 degrees, for a total of 180 degrees. If there is enough feedback signal left with a 180 degree phase shift, so the loop gain is greater than -1 at this frequency, then the loop gain becomes positive and the system oscillates all by itself -- no disturbance needed.

When you have less feedback with 180 degree phase shift than the amount needed for spontaneous oscillation, you still get resonance or "ringing" effects; after a sudden disturbance, the variables oscillate up and down for a while, gradually coming back to a steady state. The nearer you get to the critical condition, the longer the damped oscillations persist.

So in Demo 2, when you were adding perceptual lag, you were really pushing the system a little way toward the point of instability and spontaneous oscillation -- which is how the actual human control systems behave, too. It's been known for 40 years or more that there's a slight resonance in the human transfer function at about 2.5 Hz. This results in enabling the system to counteract faster disturbances than it could handle with a flat response curve, but at the expense of a slight amount of instability. Apparently, it's worth it.

Best to all, Bill P.

Date: Thu Feb 25, 1993 4:42 pm PST
Subject: place cells, PCT methodology

Bill Powers,

I need to read your reply about place cells some more, but real quick: First, about the statement that place cells represent the outputs of earlier processing make sense to me--its just like when we are doing eye movement research, the computer has to figure out, given head and eye position, where the eyes are fixating, and in addition we probably take into account the degree of eye convergence in order to compute depth. So yes, it of course represents the outputs of earlier computation.

However, I still disagree that any higher level system needs to analyze the status of the place cells. It doesn't need to go any further, (although "Place" must be somehow integrated--spatially or temporally, it doesn't matter here--with "Object), but there is no need for further processing to determine which object is further away than another. The place cells don't inform a coordinate grid--they ARE the grid. If Object X is associated with place field T54 (which again is represented differentially across cells, there is not a one to one mapping or anything like that), then I KNOW, ipso facto, Where the object x is: I experience it as such--I EXPERIENCE it THERE. No system has to say "Oh, place field T54 is active, therefore it must be there." What use would that be. Its the beginning of infinite regress--I want to stop it before it even gets started.

I need to read your post again, but it seems from your language that we are not possessing the same assumptions, but I don't know exactly what thats about. I don't think you meant the following but just in case: the location of place cells in the brain means nothing [sorry, this isn't written correctly--I am saying what I think, not the converse as I implied at the beginning of the sentence]. If they are arranged in some nice arrangment, then that's nice for neuropsychologists, but it has nothing at all to do with this conversation. I don't think you think it is relevant either but some things you said made me unsure.

Concerning consciousness, I agree and am familiar with everything you said (minus the part about the two requirements for consciousness) so I probably wasn't clear about what I meant (which I have since forgotten and will go back and look at later).

ON ANOTHER NOTE:
Last night I stayed up all night writing a research proposal for an eye movement study concerning aesthetic judgments and scene (landscape) complexity (defined in a variety of ways). Anyway, when I first sat down to write it, I thought "Now, I need to be sure that I am PCTC (PCT correct) about my methodology." So I tried to figure out how to do this--I thought about the TEST and so forth and concluded at the outset that such considerations are irrelevant to this type of study since all I want to know is where and how the eyes go and why (to put it simply). So I went ahead and began designing, and what should occur but I find myself in the end with a study which corresponds exactly to the suggestions I

made in my thesis about how to do research given a PCT perspective. To me this means that I have two independent variables (at least)--one is environmental and the other is organismic. In the study I "manipulate" the internal variable (reference state) via the instructions (and assume with good reason that the subjects will comply, and I manipulate an environmental feature which will relate directly to the organisms attempt to achieve this goal state. I could go into a lot of detail, but I am basically Expecting and Wanting interaction effects (something most psychologists seem to hope to not get).

Anyway, it was exciting to me that PCT thought is so ingrained that I would naturally design in such a way--its comforting to know that I naturally follow my own advice, even though I couldn't explicitly use the advice to make the design at the outset.

Mark Olson

Date: Fri Feb 26, 1993 9:09 am PST
Subject: Feedback Lags

[from Gary Cziko 930226.1415 GMT] Bill Powers (930225.1400) explained:

>That's the real reason for putting the lag in (aside from the
>fact that perception probably does actually involve some lags).
>The model with only a slowing factor in its output behaves too
>smoothly; it controls too well. As you increase the lag, you can
>see that the model starts to show more rapid movements. These are
>actually making its control worse, but they are making the
>difference between the model and the person smaller. So this
>tells us that we're at least in the right ballpark; the person,
>too, must contain something like the added lag.

All this helps considerably. What I am understanding is that it is not where the lags are that matters but that two act differently than one. But as a test of my understanding, I'd appreciate some quick reaction to statements I think follow from all this. Are they right?

1. If the output not quite slowed enough to prevent some overshooting with high loop gain, then adding some perceptual lag could actually make control BETTER, not worse.

2. It is theoretically (but not physically) possible to have a smootly controlling control system without any output slowing at all, just perceptual lag.

Thanks--Gary

Date: Fri Feb 26, 1993 9:13 am PST
Subject: Statistics; place cells

[From Bill Powers (930226)] Martin Taylor (930225.1530) --

>If you limit the task difficulty, the statistical limitations
>won't show up. But if you make the cursor in a tracking task

>move too fast or with too high a bandwidth, the human will lose
>control. A model with infinite bandwidth won't.

I think we're all a lot closer than before to agreeing about the legitimate uses of statistics. I have one basic quibble left, which I've actually had since about one day after I learned about the tenets of quantum physics.

Let me phrase it as a question: are there really "statistical limitations," or are there "limitations that we analyze statistically"?

My implication is that while we can analyze things like bandwidth using the interpretations and calculation methods indigenous to the world of statisticians, there are many cases in which we can get to the same answer using a different set of interpretations and calculation methods.

The basic split between Einstein and the quantum mechanicians was not, I think, over the concept of probability but over the difference between a digital and an analogue representation of the world. Einstein simply didn't believe in discrete physical variables. When you believe in discrete variables, the world consists of point-objects and point-events, created primarily by the way you characterize observations. Once the world has been characterized that way, the laws of nature reduce to logical rules rather than continuous relationships. And the only way to deal with continuous processes is to represent them as ensembles of discrete events. This naturally brings in statistics, because statistics has to do with discrete events that are counted, not continuous variables that are measured. Statistics deals with qualitative events.

It's always seemed strange to me that at the foundations of quantum mechanics we have Schroedinger's Equation, which is cast in terms of continuous variables. It seemed to me from the start that if quantum effects were truly quantized, then the mathematics with which we analyze them should be integer mathematics, not an equation in continuous variables using partial derivatives. The moment you allow partial derivatives into the picture, you're relying on the existence of an underlying continuum. It has always seemed to me that quantum mechanics rests on a trick -- pay no attention to that man solving a continuous equation behind the curtain!

This same continuous-discrete distinction keeps showing up in conversations on CSG-L. There seems to be a strong desire to reduce continuous quantitative relationships and processes to discrete happenings. This shows up sometimes in the form of puzzles that are puzzling only because of using qualitative descriptions instead of quantitative ones. I mentioned one example yesterday: if control systems correct error, but their actions are based on error, then if the action corrects the error there is no error to drive the action and the action must cease. which permits the error to return ... . On the basis of that kind of analysis, we would expect all actions to consist of oscillations: action, no-action, action, no-action....

If one is committed enough to such a view of the situation, rationalizations ("renormalizations") can always be found. Perhaps action really consists of virtual pairs of actions and anti-actions that are continually forming and annihilating each other, with a statistical bias on the preponderance of one over the other that gives, on a macroscopic scale, the appearance of a continuous amount of action between the actual stable states. What we need then is a statistical law that will predict the average state of the action-no-action antiparticles, so that we can predict the intermediate states that we observe on the macroscopic scale.

Of course what we end up with is an extremely complex way of computing the state of a system, the same state we could predict, on the basis of assuming continuous relationships, with a couple of lines of algebra.

If I were asked to predict the behavior of a person in a tracking experiment when subject to disturbances outside the bandwidth of the control system, I would not use statistics but differential equations, or better yet simulate the situation and let a model predict the result. I don't know of any statistical method available now that could even get me close to a correct prediction of the person's handle movements.

Of course we could say that IN PRINCIPLE we could calculate the dynamic responses of the person's input functions, comparators, output functions, and muscles from basic physical principles. Given the relationship between stimulus intensity and neural impulse rate, and the distribution of impulse rates around the average rate for a given stimulus, and the forms of the neurochemical functions in the neurones, and so on at considerable length, we could derive the overall system response and predict the behavior for a given disturbance wave form. We could do this in principle, if we were God.

As we are not God, all we can actually do is make up some assumptions that permit us to apply a statistical analysis, adjusting the assumptions until the analysis gives the right answer. I don't see any advantage in doing that, when there are far simpler methods available that give better answers than we have yet obtained from any more complex approach.

>>As I look at Bill's description of the compensating system, it
>>seems very like the imagination loop.

>On re-reading, I realize this is subject to misinterpretation.
>What I mean is that the function R looks like, and has much the
>function of, the model that we seem to be developing as a
>necessary component of the imagination loop.

Glad you said that. Whether we're talking about Ashby's system or the imagination loop in a hypothetical control system, there's a hidden control system needed to make the scheme work. It's the one that looks at the outcome, the essential variable in Ashby's case and the model's behavior in our case, and compares the actual result with the desired result. If the desired result is wrong, the R function or the model has to be adjusted to correct the error. SOMEBODY or SOMETHING has to know the state of the outcome in order for the correct R function to come into existence, or in order for the right model to be constructed. It's all very well to say "given the perfect regulator R" or "given a correct imaginary model M", but the nitty-gritty modeling problem is WHO GIVES IT?

Oh, I forget: evolution gives it. Whew, I thought we had a problem there for a minute.

Mark Olson (930225.1737) --

How about putting that information at the start of your posts, like

[From Mark Olson (930225.1737)]

The date and time in the header (which I used above) aren't necessarily the same as the data and time of your post -- and anyway it's nice to know who's talking without having to scroll down to the end of the post to find the signature.

>I still disagree that any higher level system needs to analyze
>the status of the place cells.

I agree that there's no need for analysis in order for a perception of placeness to exist. The signal coming out of a place cell, as you say, IS that perception. No problem there.

The problem arises if you don't just say "OK, there's a place signal, next question."

I say to you "the keys are on the hall table -- would you hang them on the hook by the refrigerator?" Now we have several placeness signals -- the place where the hall table is, and the place where the hook is. You have to remove an object from one place and put it in another place. Assuming most of the action as given, one problem that remains is which way to go with the keys after you locate them in the hall-table place.

In short, something has to consider the two places (hall table and hook beside refrigerator) and perceive the DIRECTION of a path between them. "Direction" is not a place signal. It is a signal derived from several place signals, by a higher level of system. The perception obtained by the higher system is not a place perception, but a perception of direction (or transition if you like). The direction in which you walk can't be controlled unless you perceive the direction in which you're moving, and you can't select a reference direction until you perceive the RELATIONSHIP (another level still) between where you are and where you want to be.

As to the physical relationship between place cells and the meaning of the place signals, we understand each other. The physical arrangement is irrelevant. I was only pointing out that "mapping" in itself is unimportant, and that only the connectivity of the cells matters. The mapping effect might be a consequence of the system evolving to make spatial computations more efficient; shortening the computing pathways. But there is nothing internally that can appreciate spatial relationships by literally looking at the map.

That was a nice story about the research proposal. Knowing PCT has indeed distorted those nice neat stimulus-response minds we inherited.

Best to all, Bill P.

Date: Fri Feb 26, 1993 10:04 am PST
Subject: thinking letters

[From Rick Marken (930226.0900)] Bill Powers (930225.1400)--

>My impression was that the EEG waves were run through a pattern
>generator that would type out the letter, with the person

>controlling the letter typed just by imagining it.

This would be quite impressive, indeed, especially if the person just pictured a sequence of letters in imagination (say, 10 different letters pictured one after the other at the rate of about one letter every 2 seconds). If all (or 90%) of the letters were detected correctly in the EEG then I would be very impressed; I would be amazed if one could detect in the EEG the neural currents that correspond to imagined letter perceptions. But, if they can be detected, then they can; it would be a nice, non-messy way to test the cortical basis of the PCT hierarchy.

> Would someone who has actually read the article or seen the demonstration tell
>us what this is really about?

I agree;a detailed description of the study would really help.

Best Rick

Date: Fri Feb 26, 1993 12:44 pm PST
Subject: Attracting Psycholoquists

[from Gary Cziko 930226.1952 GMT]

There was some discussion a while back about a proposal made by Rick Marken to attract Psycholoquy (a moderated e-mail list of "psycholoquists"--rhymes with "ventriloquists" and some of us know why) subscribers to CSGnet. Rick proposed submitting an add something like this for posting on that network:

"I would like to put a notice in Psycholoquy indicating the availability of an Internet discussion group on purposeful behavior. The group, csg-l, discusses the implications of a control model of purposeful behavior for studies of language, information processing, AI, chaotic attractors, social behavior, neural networks, psychopathology, evolution, operant behavior, consciousness, etc. You can subscribe to csg-l by ..."

About three people thought this was a good idea. About another three expressed reservations. I don't recall receiving reactions from the two people (other than me) who would be most "impacted" by a sudden influex of new CSGnetters: Bill Powers (since for some reason he seems compelled to answer every question, stupid or otherwise, put on the net--including mine); and Greg Williams (who puts together _Closed Loop_).

Before proceeding with this, I would like to get the reaction of other CSGnetters who have not yet given their opinion--especially Bill and Greg.

--Gary

Date: Fri Feb 26, 1993 12:46 pm PST
Subject: Re: Feedback Lags

[Martin Taylor 930226 15:00] (Gary Cziko 930226.1415)

Let me see if I understand.

Slowing factors incorporate part of past output into present output, so that an output spike comes out somewhat smoothed. Lags delay, so that an output spike is still a spike, but later than it might otherwise have had its effect.

In a discrete system, slowing factors make the system ergodic, whereas lags do not.

Hope this is reasonable and not totally misunderstood.

Martin

Date: Fri Feb 26, 1993 1:30 pm PST
Subject: Re: Statistics; place cells

[Martin Taylor 930226 15:10] Bill Powers 930226

I think one of the things that causes you difficulty with statistics is some kind of linkage between statistics and discrete categorization. There isn't any issue as to whether the statistics deal with continuous or with discrete systems.

As I remember mt undergrad quantum theory, the discreteness falls out directly from the Schroedinger equation. In verbal intuitive terms, only at certain discrete values of the variables do the virtual orbits reinforce (recur), so that the probability densities for the real object (electron) sum to non-zero velues. The discrete is a natural consequence of the continuous, and the same applies in the matrix representations (which we have applied to the activities in recurrent neural networks).

I don't think I have been talking about any of this when I have been
talking about statistics.

\> Whether we're talking about Ashby's system or
\>the imagination loop in a hypothetical control system, there's a
\>hidden control system needed to make the scheme work. It's the
\>one that looks at the outcome, the essential variable in Ashby's
\>case and the model's behavior in our case, and compares the
\>actual result with the desired result. If the desired result is
\>wrong, the R function or the model has to be adjusted to correct
\>the error. SOMEBODY or SOMETHING has to know the state of the
\>outcome in order for the correct R function to come into
\>existence, or in order for the right model to be constructed.
\>>It's all very well to say "given the perfect regulator R" or
\>"given a correct imaginary model M", but the nitty-gritty
\>modeling problem is WHO GIVES IT?
\>
\>Oh, I forget: evolution gives it. Whew, I thought we had a
\>problem there for a minute.

Yes, of course one must have a real control system in the development of R. (Evolution could do it, too, if the world is stable enough. So could reorganization of a kind, if the world isn't stable enough for evolution but is stable enough within a lifetime). In the PCT imagination loop (not the Ashby formulation, perhaps -- I still haven't read him), R works with and is modified through a normal ECS.

(Bill Powers 930225.1400)

> The impression I get is that you consider noise to be a rather large
>component of the error signal, especially at zero error.

Well, at zero average error, the noise is all there is, so indeed it is a large component of the error signal.

>When I hold an arm out straight and try to keep the forefinger
>still, the finger wobbles laterally around by no more than 2 mm
>(for a while anyway). This is 1/400 of the total range. The
>question is, how are we to speak of that 2mm of noise in the
>control system? If we compare the average error, which is less
>than 2 mm, with the RMS wobbles, then it will seem that the error
>signal consists MOSTLY of noise, with very little signal. My
>impression is that this is how you have been thinking about it.

You are talking about output here. To me, that isn't very relevant. It does put an upper bound on how bad things can be (the perceptual noise presumably can't be greater than one might deduce from the output wobble), but that's not very interesting. The interesting point is how fast one can perceive to whatever level of accuracy is important.

>If the random fluctions in error signal were still significant in
>comparison with the average magnitude of the error signal, we
>would expect to see significant fluctuations in the force
>opposing the disturbance. But we don't. The force is essentially
>steady.

As I mentioned a few posts ago, and you quoted, it is unlikely that we would develop control systems that tracked the noise fluctuations, or that we would (easily) become conscious of the noise in the system. What evolutionary purpose would be served? We need to perceive and to act in respect of things in the outer world. Anything else is at best wasted effort, and at worst dangerous. But we need to perceive as fast and accurately as we can, and that is noise-limited in lots of places other than just the neural firings (we really don't know jsut where the information-carrying aspects of neural signals are--I suspect that much more than simple inter-spike intervals are effective). You pay for your output steadiness in output speed (or lack of speed).

It all comes down to: what we can't perceive, we can't control. You get information at a limited rate about any perceptual signal, and the higher the perceptual level, the lower (probably) is the rate.

>I think you may have been doing the modeling equivalent of
>dividing by zero. By selecting conditions under which the average
>error is zero, you artificially inflate the importance of random
>noise in the error signal. Under other conditions, the error can
>depart from zero by many times the magnitude of the noise
>fluctuations (still remaining reasonable small in terms of the
>whole range of control). Suddenly the signal-to-noise ratio of

>the error signal looks high instead of essentially zero.

I really don't distinguish between situations where the error signal is near zero and where it isn't. And the signal-to-noise ratio isn't much of an issue either, though it kind of comes out in the wash as a surrogate. If you get information about the position of a target at 1 bit per second, you can halve your uncertainty of perception every second, and this limits how fast you can bring your finger to point accurately at it. If you get information at 10 bits/second, you can halve your error 10 times per second, and could in principle get to an RMS error of 1/1000 of full range in one second (given a few constant factors). Schouten's data suggest that one can discriminate between two lights at around 140 bits/sec under the conditions of his experiment, so we may judge that the limiting factors in pointing a finger at a light that comes on abruptly are not those of the visual perception of the light. Other perceptions (kinaesthetic?) are more limiting.

How accurately or stably you finally point is interesting only in providing an upper bound to how noisy the residual percept might be. (I have notions on that, too, that involve dead-zone perceptual signals, but it is obviously too soon to introduce them into the discussion).

Martin

Date: Fri Feb 26, 1993 1:39 pm PST
Subject: Lags vs. Slowing

[from Gary Cziko 930226.2100 GMT] Martin Taylor 930226 15:00

>Slowing factors incorporate part of past output into present output, so
>that an output spike comes out somewhat smoothed. Lags delay, so that
>an output spike is still a spike, but later than it might otherwise have
>had its effect.

I had been using lag and slowing interchangeably, but I like your distinction and so will use if from now on unless somebody who knows says it's wrong.

Maintaining this distinction, in Demo2 Bill Powers uses slowing (not lags) on both the output (action) and input (perceptual sides). Maybe it should be otherwise?

>In a discrete system, slowing factors make the system ergodic, whereas
>lags do not.

I'm not sure I understand this. I even looked up "ergodic" to make sure it hadn't changed its meaning from the last time I looked it up.

From the OED2 (which I have on-line):

>ergodic (3:r'godIk), a. Math. Of a trajectory in a confined portion of space:
>having the property that in the limit all points of the space will be included
>in the trajectory with equal frequency. Of a stochastic process: having the
>property that the probability of any state can be estimated from a single
>sufficiently extensive realization, independently of initial conditions;
>statistically stationary. Also, of or pertaining to this property.

Do you mean slowing "smears" things out while lags keep things jumping around without adding any smearing?--Gary

Date: Fri Feb 26, 1993 2:20 pm PST
Subject: Transport lags and slowing

[From Bill Powers (930226.1300)] Gary Cziko (930226.1415) --

I guess I've misled your intuition a bit. Let's see if I can give a short answer to your comments (I deleted a 2000-word answer, so be grateful). I have to take the second point first.

2. It is theoretically (but not physically) possible to have a
>smoothly controlling control system without any output slowing
>at all, just perceptual lag.

This is easiest to handle in frequency terms. A control system will be unstable if its loop gain is greater than unity at a frequency where the cumulative phase shift around the loop is 180 degrees. That phase shift converts negative feedback to positive for any signal at that frequency. Noise will then suffice to set off spontaneous oscillations.

A pure transport lag delays a signal without changing its amplitude. For a transport lag of t seconds, the frequency that suffers a 180-degree lag is 1/(2t). So a transport lag of 1 second would put a 180 degree phase shift into a frequency of 1/2 Hz, and a lag of 1 millisecond would cause a lag of 180 degrees in a 500 Hz sine-wave signal. If you put a transport lag into a stable system that has no other lags at all, then no matter how short the lag t, the system will oscillate at the frequency 1/(2t). So a theoretical control system with a pure transport lag in it (anywhere in the loop) and a loop gain greater than -1 will oscillate like mad no matter how short the lag is. You can't stabilize a control system by adding a transport lag: quite the opposite.

Now your first point.

>1. If the output were not quite slowed enough to prevent some
>overshooting with high loop gain, then adding some perceptual
>lag could actually make control BETTER, not worse.

Our basic control system has a (leaky, but that's irrelevant here) integrator in its output. The reason for using a leaky integrator is that if we just made the output proportional to error, the loop would oscillate at the frequency of iteration of the program. There is an irreducible 1-iteration transport lag in the digital computation, and with a loop gain greater than unity the system would oscillate at a frequency of 1/(2t): one iteration up, the next down.

In order to prevent this oscillation, we have to slow the changes in output so that on any one iteration, the error can't be overcorrected. That's what the integrator does. If there happens to be a transport lag of several iterations, deliberately introduced into the program as we do in Demo 2, then even more slowing is required; the error should not be overcorrected in however many iterations are represented by the transport lag.

So your intuition went at the problem backward. We would not need any slowing at all if it weren't for the transport lag. The longer the transport lag, the more slowing is needed to prevent oscillations.

When you solve the control equations algebraically, ignoring time, you are in effect working with a system containing zero transport lags. This allows you to treat the component equations simultaneously, which literally means at the same instant of time. All real systems, however, have lags in them, and all computer simulations have a lag of at least one iteration.

I may not have made my main point clearly enough. Our problem is not to get the control system with a single leaky integrator at its output to control as well as a person does. By using a smaller slowing factor and a larger gain, you can make the model control so well that the "cursor" trace is a straight line on the screen -- far better than the person can do. By increasing the slowing and lowering the gain, you can make the model control worse -- and at the same time, control more like the real person. At some point the increase of slowing starts making the match to the person worse again; that minimum RMS difference between model and person is the best you can do with only these adjustments. You can make the model control even worse by making it still slower, but now it will be controlling differently from the way the person does, more sluggishly.

In the second part of the demo, a perceptual lag is introduced. This makes the model control worse in a new way: by overreacting to fast changes in the disturbance instead of acting too slowly. The perceptual lag partly destabilizes the model, the lag increasing the phase shift. This increases the higher-frequency wobbles of the handle, in just the way needed to fit the real person's wobbles better. In fact, the RMS difference between the model's behavior and that of the person is reduced by 30 to 50 percent by adding the correct amount of lag.

The model is not made to control as well as possible, but only as well as the person does. If it couldn't control better than the person does, we wouldn't be able to find the parameters that fit the model to the person by making the model control worse.

Hope that makes it clearer.

Best, Bill P.

Date: Fri Feb 26, 1993 2:46 pm PST
Subject: place and eye movements

[From Mark Olson 930226 4:00]

Ok, I'll try to remember to put in the above stuff that is in brackets--doesn't everyone receive a header though, with the same information? When replying to a post, I would put the same info in about which I am responding to, but its difficult since I read net stuff from one system and mail from another. This should change soon however, cause I just acquired a real user-friendly program for IBMs that allows me to construct messages off line and then send them, so I don't feel like I'm taking a timed test while I write a message (a timed test which may end at any random moment when someone decides to call me).

Anyway

Date: Fri Feb 26, 1993 3:36 pm PST
Subject: place and eye movements

From Mark Olson (same time as previous message)

Oh, I guess I should say [930226 4:30]

Sorry that the last message got cut off.

Could someone make a PCT prediction for me concerning the following? Evidently, if a subject fixates on a point and is told to move to a point to the left (which they can see but is not in foveal vision), they will saccade directly to that place where the second dot is. If the second dot is moved during the saccade, they will land where they "should" have given the original position, and then adjust accordingly. After some time, the saccades will adjust so that they do land on where the dot has been moved to (its always moved to the same place).

This indicates to me that an alteration has occurred such that the representation of Where the dot is is mapped onto a different reference for spatial location that the motor system is going to use to determine the final location of the saccade. That's right, so far, isn't it--position p44 no longer maps to reference setting 91 but now reference setting 77, for instance.

Ok, now, given that, what would we predict would happen if after fully altering these relationships, the subject goes to reach for an object. Will he (A) miss the object initially because the visual system isn't calibrated with the motor system, or (B) not miss the object because visual feedback will indicate that he is undershooting or overshooting? I could see it either way.

Related to this, I am told that if a person is told to fixate on a letter on a screen (say G) and the person receives visual feedback concerning where his eye is fixating by having the computer place a crossbar there, then the person will move his eyes in order to get the crossbar on top of the G. I don't think its that they are "trying harder to look at the G" but rather that they are "trying whatever works to get the crossbars over the letter. (For this reason, subjects are not given such feedback during an experiment).

I think this is a fascinating type of feedback, thought I cannot at the moment explicate why it is so unique. Anyway, if we altered the offset such that the feedback was innaccurate, what would happen when the subject gets up to leave? Will he run into the doorframe or something similar. I know the effect wouldn't last long, but would there be a notable effect and for how long (one second, one minute?)

These are questions that I would imagine Bill Powers is familiar with from designing Little Man.

Mark

Date: Fri Feb 26, 1993 3:38 pm PST

Subject: Psycholoquy; statistics

[From Bill Powers (930226.1500)] Gary Cziko (930226.1952 GMT) --

I've been thinking over the ad for Psycholoquy. My hesitation comes partly from the idea of getting a lot of new people on and having to start again from scratch. On the other hand, why not? So I guess I've come down in favor of the idea.

P.S. In the Demo 2 experiment, part 2, the perceptual lag is
added as a true transport lag, not a second slowing factor.

Matrin Taylor (930226.1510) --

>As I remember my undergrad quantum theory, the discreteness
>falls out directly from the Schroedinger equation. In verbal
>intuitive terms, only at certain discrete values of the
>variables do the virtual orbits reinforce (recur), so that the
>probability densities for the real object (electron) sum to non-
>zero values.

That's interesting. I should think it would follow that near these discrete values you would have orbits that nearly recur, so the center of the value would be surrounded by lifetimes that were long but not permanent. It sounds pretty continuous to me.

I've often thought that there must be a family relationship between quantum phenomena and macro phenomena like standing waves.

>In the PCT imagination loop (not the Ashby formulation, perhaps -- I still
>haven't read him), R works with and is modified through a normal ECS.

The model would work with a normal ECS, but a separate process is needed to compare the output of the model with the way that output "ought" to look, and alter the model to make the difference smaller. Once the model is correct it is used as part of an ECS, but before it is correct a very different sort of control process has to be at work.

RE: noise and error signals.

>You are talking about output here. To me, that isn't very relevant.

The output is directly driven by the error signal. With some allowances for filtering by the masses of the arm, the wobbles in the output reflect deviations of the error signal from a constant frequency. There may be some other sources of noise after the error signal, but the error signal is greatly amplified, so subsequent noise sources would have relatively less effect.

>But we need to perceive as fast and accurately as we can, and
>that is noise-limited in lots of places other than just the
>neural firings (we really don't know jsut where the information-
>carrying aspects of neural signals are--I suspect that much more
>than simple inter-spike intervals are effective).

I don't know of any other way for neural signals to carry information, considering how they work postsynaptically. Or to put it differently, so far I haven't seen any other mode of neural transmission that looks feasible considering the nature of the receiver.

>If you get information about the position of a target at 1 bit per second,
>you can halve your uncertainty of perception every second, and this limits
>how fast you can bring your finger to point accurately at it.

I would put this the other way around. If the physical properties of the system result in halving your error every second, then it can be calculated that "information" must be flowing at the rate of one bit per second. The rate at which error is corrected can be found from a model of the physical processes involved, without using the calculation of information rate. The information rate is a derived quantity, not a causal factor. It does not explain why the system behaves as it does. It isn't the information rate that limits how fast you can point your finger; it's how fast you can point your finger that determines what information rate you will calculate.

The way you put it, it would seem that if all you know about the system is that it is getting information at a rate of 1 bit per second, you can predict that it will halve its error every second. That doesn't follow. FIRST you have to know how rapidly the system will reduce its error, on physical grounds. Only then can you estimate the limiting equivalent information rate. Information isn't a substance that gets passed around; it's a way of characterizing physical processes after the fact.

I expect that we're going to be arguing about this for some time to come.

Best to all, Bill P.

Date: Fri Feb 26, 1993 3:43 pm PST
Subject: Re: Ashby

[From Allan Randall (930226.1730)] Bill Powers (930224.1730) writes:

>>Ashby is saying that error
>>control cannot be perfect since it relies on detecting errors.
>
>...He seemed to be claiming that a compensatory system could prevent an
>error in bath temperature by seeing that someone is coming with a cold flask,
>and acting NOW.

You keep referring to the "compensatory regulators" (your term, I think?) as if they are the main point of Ashby's book. In fact, Ashby makes it clear that error control is much more important.

<Ashby said
<<... the thermostat-controlled water bath, which is unable to say "I see someone
<<coming with a cold flask that is to be immersed in me -- I must act now."
<
<I assumed this to imply that there IS a kind of system that could
<achieve control by anticipating a disturbance of this sort.

All of your arguments against Ashby are prefaced with phrases like "I took this to be an implication that...," "...is used to imply that...," "Note the implication that...," "...the unspoken assumption that...," "I assumed this to imply that..." You are reading heavily between the lines here, and ignoring what Ashby actually *says* about error control, that it is "of the greatest importance and widest applicability. The remainder of this book will be devoted to it. (The other cases [compensatory systems] are essentially simpler and do not need so much consideration.)"

You say:
<Ashby definitely is assuming that a [compensatory] regulating
<scheme like the one above could actually achieve significant
<regulation. It could not, in the real world.

Why "definitely"? I still do not see Ashby making the above claim. Again, I would like to see specific references in the book to where he *explicitly* makes this claim. None of the references you give actually state that compensatory systems are superior in the real world to error-based systems.

Incidentally, I have no problem with your criticisms of the compensatory systems. I would be the first to toss Ashby in the trashbin if I thought he was really saying what you claim he is. To even think of some of these diagrams representing the kinds of control systems responsible for cognition is laughable, and utterly indefensible.

<The problem is that his is a qualitative statement: either
<regulation is perfect or it isn't.

Why is he not allowed to talk in those terms? He's talking about *perfection* here, not usefulness. Either it is or it isn't perfect. Error control isn't. Mathematically, error control is imperfect by definition. This is what Ashby meant when he said it was a *fundamental* property.

<If, on the other hand, he had considered both kinds of system
<quantitatively, the judgment would have been that while control
<by error and control by compensation are both necessarily
<imperfect, control by compensation is inherently far more imperfect.

I think this is incorrect. Compensatory control, unlike error control, has no such inherent mathematical limitation. This is an abstract mathematical statement, not a nuts-and-bolts issue. Perfect compensatory systems will only work for uninteresting, simple environments, like the simple tables of small integers that Ashby uses. Nonetheless, in principle they can be perfect. Error control systems, on the other hand, are inherently imperfect, and much more powerful.

>>The fact that it has detected error at all means that its
>>control *cannot* be perfect.
>
>No, it cannot be perfect. However, it can be so close to perfect
>that no human eye could tell the difference.

Then you agree with Ashby. Error control cannot be perfect.

>>The Law of Requisite Variety (or information) says simply that
>>the complexity in the disturbance must be made informationally
>>at least as low as the output channel capacity if the control
>>system is going to successfully control against this disturbance.
>
>But this is just a complicated and abstract way of saying that
>the output of the control system must quantitatively oppose the
>effects of the disturbance.
>...
>This is one of those true but useless statements.

It is "complicated" and "abstract" only because it is not in your language. You do not like information theory. Fine. But the fact that Ashby thinks in these terms does not make him a "bad guy" fit for the Devil's Bibliography.

<My biggest guffaw at Ashby's reasoning came on p. 223-224, where
<he goes through the logical error that practically every neophyte
<in control theory goes through. "Control systems correct error.
<But their actions are driven by error, so if there's no error,
<they can't act, and if they can't act, they can't correct error."

Therefore, they cannot control perfectly. You yourself have already agreed with this point. I don't see how you can reasonably argue otherwise.

>>Again, I'd really like to know where in the book he says this.

>I've already noted the sections, but here they are put together:

None of these sections make the claims you attribute to Ashby.

>...The arguments that follow [Ashby's] are based on the idea that
>error-controlled system are inherently -- and significantly -- imperfect.

It is the "significantly" part I'd like to see references for.

<The only way for R to be a perfect regulator is for the variables
<it handles to be logical or binary and the calculations involving
<them error-free. The examples Ashby uses always use small whole
<numbers or logical propositions, in which small errors never
<occur. A disturbance of 3 units is canceled by a regulator output
<of -3 units: perfect regulation. If we considered real
<regulators, even those involving binary variables, we would find
<nonzero errors, and the in-principle perfection would vanish.

Then you would agree with Ashby when he says of error control:

"Thus the presence of continuity makes possible a regulation that, though not perfect, is of the greatest practical importance. Small errors are allowed to occur; then, by giving their information to R, they make possible a regulation against great errors. This is the basic theory, in terms of communication, of the simple feedback regulator." p. 224

I can accept (for now) that you dislike the "in terms of communication" part, but the "simple feedback regulator" part is essentially correct.

<...I can't find the reference right now, but at one point
<Ashby "proved" that the maximum loop gain of a stable negative
<feedback system (implemented in discrete operations) is less than -1.

If you can find this, I'd be interested. (I'm not disagreeing with you on this one - just interested).

>>The compensatory regulators I see in Ashby are mathematical
>>idealizations... a route for Ashby to understand
>>the real issue: error control.
>
>But when Ashby finally got to error control, where did he go with
>it? To "The Markovian Machine," another exercise in manipulating
>discrete variables and small integers...
>...
>Ashby was actually an idol of mine after I read "Design for a
>brain" shortly after it was published. He said many true and
>useful things about control and negative feedback in that book,

I guess the higher you place 'em the harder they fall. Ashby makes it clear in the introduction that the book is from a very different perspective than "Design for a brain." He urges the reader not to read the book in isolation, but to consider the two books as complimentary. "An Introduction to Cybernetics" is meant to be an abstract, mathematical treatment of control theory from the point of view of information theory. I really think you are doing Ashby an injustice here. I can accept that you do not appreciate information theory, but I would not place you in the "Bad Guy" camp because of this. Please give Ashby the credit he is due.

Allan Randall, randall@dciem.dciem.dnd.ca
NTT Systems, Inc. Toronto, ON

Date: Sat Feb 27, 1993 11:35 am PST
Subject: Re: Lags vs. Slowing

[Martin Taylor 930227 14:20]
(Gary Cziko 930226.2100)

>>In a discrete system, slowing factors make the system ergodic, whereas
>>lags do not.
>
>I'm not sure I understand this. I even looked up "ergodic" to make sure it
>hadn't changed its meaning from the last time I looked it up.
>
>>From the OED2 (which I have on-line):
>
>>ergodic (3:r'godIk), a. Math. Of a trajectory in a confined portion of space:
>>having the property that in the limit all points of the space will be included
>>in the trajectory with equal frequency. Of a stochastic process: having the
>>property that the probability of any state can be estimated from a single
>>sufficiently extensive realization, independently of initial conditions;
>>statistically stationary. Also, of or pertaining to this property.
>
>Do you mean slowing "smears" things out while lags keep things jumping

>around without adding any smearing?--Gary

Basically, the definition is what I meant, and your interpretation is a reasonable one. The way I look at it, in a discrete system without slowing, whatever happens at time t0 at the output is reflected after a transport lag in some probability distribution of effects at the sensory input at time t0+delta. Assuming that the world does not "smear" the effect over time, the output at t1 later than t0 but earlier than t0+delta will incorporate nothing of the t0 effects, and the sensory input at t1+delta will be independent of the sensory input at t0+delta. The sensory input that was affected by the t0 output will affect the output again at t0+loop delta, and that at t1 will affect the output at t1+loop delta. The two sets of signals around the loop will forever remain independent. If the world reacted differently to each of the outputs, it is quite possible that the dynamics of the system represent two (or more) completely independent waveforms with different probability distributions for the states.

If you put in a smoothing factor, whatever happens at t0 eventually affects the signals that circulate from all starting times, so that there is no privileged moment in time. You can compute the state probability distributions no matter when and with what value you start.

I think all these different kinds of statement are just ways of putting the same thing. Martin
Date: Sat Feb 27, 1993 11:50 am PST
Subject: Re: Psycholoquy; statistics

[Martin Taylor 930227 14:40] Bill Powers 930226.1500

>Matrin Taylor (930226.1510) --
>
>>As I remember my undergrad quantum theory, the discreteness
>>falls out directly from the Schroedinger equation. In verbal
>>intuitive terms, only at certain discrete values of the
>>variables do the virtual orbits reinforce (recur), so that the
>>probability densities for the real object (electron) sum to non-
>>zero values.
>
>That's interesting. I should think it would follow that near
>these discrete values you would have orbits that nearly recur, so
>the center of the value would be surrounded by lifetimes that
>were long but not permanent. It sounds pretty continuous to me.
>
>I've often thought that there must be a family relationship
>between quantum phenomena and macro phenomena like standing waves.

Cute, and I think correct. I think you are probably talking about the lifetimes of unstable particles, but I don't know enough quantum mechanics to be sure. When you say "long" I wonder how times 10^^-36 seconds you mean. I think 10^^-18 sec would be enormously long here. But that's way beyond my real understanding. I was only reaching back a few decades a plucking something out of a fading memory.
----
on statistics:

>I expect that we're going to be arguing about this for some time to come.

Perhaps. But is it worth making the rest of the CSG-L readers put up with the argument now, or should we let concepts percolate and perhaps arrive at a condition in which some new way of putting it will help us see things the same way (if possible)?

I'm prepared to let it drop for now. I don't disagree with much in your last post, except its relevance to the points I am trying to make.

Martin

Date: Sat Feb 27, 1993 1:00 pm PST
Subject: let's play e-coli!!

[i.n.kurtzer(930227.1430)]

i am up to my neck in the crap they call education here in the piney woods of east texas so i'll be brief. i thought of a game where someone could see what it would be like to be an e-coli (put fat quotes around that) . for hannukah, years ago, my parents would hide presents in the house and i would have to find them. but my search was aided by ema and aba (that's hebrew for ma & pa right?,whatever) saying either "hot" or "cold"; so whenever i was getting closer to the hidden present they would say "hot" and if i was getting further away from the present they would say "cold". eventually, i found my presents; heck, i bet i could even do it with my eyes closed. anyway, i don't know how far the analogy goes especially considering how the loop through is completed through someone else. i said it was like the e-coli because at any moment it could go from hot to cold then i would keep changing direction until i was back on track. it's a cute game.
----
to bourbon: are you planning to go to SWPA this year; it will be in Corpus Cristi and i would turn cartwheels to ask a few questions to Bib again, heh heh. also, psychi wants you to give a speech here on pct (i heard through the grapevine).
----
to greg williams: i'll be "finished" with my sparta paper at the end of march and will distribute a few to get some possible problems straightened aout,afterwards i'll send you the final copy.

are my lines correct yet? i.n.kurtzer

Date: Sat Feb 27, 1993 3:09 pm PST
Subject: psycholoquy advertisement

[Avery Andrews 930301.03930]

I'm not seriously opposed to the advertisement, in spite of what I've said about it, but I would recommend that it be light on the revolutionary claims, and focus more heavily on the idea that feedback control has been poorly conceptualized and underexploited. I think it might also be nice to get the Devil's Bib. material that people have contributed so far organized. One thing that should be done is that people's contributions should be checked by at least one and preferably 2 people, to eliminate or qualify idiosyncratic readings & misinterpretations. I guess this means that somebody else should get into the Ashby debate, tho my hands are full at the moment.

Avery.Andrews@anu.edu.au

Date: Sat Feb 27, 1993 8:25 pm PST
Subject: Psycholoquy; Slowing

[from Gary Cziko 930227.1515 GMT]

Bill Powers (930226.1500) says:

>I've been thinking over the ad for Psycholoquy. My hesitation
>comes partly from the idea of getting a lot of new people on and
>having to start again from scratch.

Thinking this way I would never teach the same course twice.

>On the other hand, why not?

Exactly

>So I guess I've come down in favor of the idea.

Good. Now what does our archivist think about this? (Greg, are you still out there? Or have the recent midwest snows buried you and your computer?)

>P.S. In the Demo 2 experiment, part 2, the perceptual lag is
>added as a true transport lag, not a second slowing factor.

Now I'm confused some more. Here is what you say on the relevant screen of Demo2:

==
This time, after the experimental tracking run, the model will include a sensory lag adjustment, as well as the adjustments you've already seen. This is accomplished by using a slowing factor just as in the output function . . .

A sensory slowing factor of 5 means that after a sudden change in cursor position, the perception of the cursor position changes by 20 per cent in each 1/30 second interval--20 percent of the remaining distance to the final value. This corresponds to a sensory time constant of about 0.15 seocnd, a reasonable value for visual perception.
===

Can you see why I'm confused. Sounds like you're describing the same type of leaky integrator as for the output side. If it is just transport lag, then what exactly does the slowing factor here refer to?

Now that I've thought about this some more, it seems to make sense. You use the leaky integrator on the output side since that is supposed to model real physical movement--things can't move instantly and any peaks in the output signal get smoothed out by stuff like mass (the heavier my arm, the more slowing, right?). On the perceptual side we are dealing with only sensory transduction and neural currents. This takes time, but there is not any acceleration or deceleration of masses (except for little ones like eardrums). So the perceptual signal is not changed by the system, only delayed. Delay can only be "bad", but the leaky integrator can be good since it smooths out actions which with high negative loop gain would otherwise overshoot the mark and lead to instability.--Gary

Date: Sun Feb 28, 1993 8:41 am PST
Subject: need cognitive theory references

from Ed Ford (930228:0920) To Rick, Gary, et al....

A close friend needs good references in the current literature (books and articles) on the best explanation of cognitive theory. I would appreciate anything you could offer.

To All..

In two hours, am headed for the cold climate of Green River, Wyoming, to work for a week in a school district, teaching my program Teaching Responsible Thinking (including PCT) to administrators, counselors, teachers and parents. They in turn will teach school district parents needing helping in dealing with their children.

Best, Ed.

Date: Sun Feb 28, 1993 9:03 am PST
Subject: OOOPS on Demo2; Where's Greg? Hot & cold

[From Bill Powers (930228.0930 MST)] Gary Cziko (930227.1515) --

Omigod, you're right. I just looked at the source code for Demo 2, and the perceptual lag is indeed a slowing-factor lag, not a transport lag. I'm glad the writeup was correct. I think I meant to change it to a transport lag and just never actually did it. Thanks for the correction.

However, the remarks about transport lags in recent posts still hold true.

Yeah, where is Greg? The last I heard, his son Cam had a bad bone break and was in the hospital. What's up, Greg?

Isaac Kurtzer(930227.1430) --

The "hot-and-cold" game is indeed just like the E. coli experiment!

Your line lengths are just fine now.

Best to all, Bill P.

Date: Sun Feb 28, 1993 5:51 pm PST
Subject: What's up with Greg.

From Greg Williams (930228)

What's up with me? Same as ever, of course.

No problem with the ad, here. Let the new netters arrive -- I'm ready. But I hope that some more OLD netters will agree to let me use their posts in future CLOSED LOOPs (I need copyright permission). You (might) know who you are....

The doctor decided not to open up Cam's arm -- just stomped on it a bunch in the operating room. All going well now.

Did anyone like or not like the last CLOSED LOOP? (Hint, hint.)

As (really!) ever, Greg

Date: Sun Feb 28, 1993 5:57 pm PST
Subject: P.S. to Isaac K.

From Greg Williams (930228 again!)

Thanks in advance for the paper-to-come. And your lines are lovely. I think I can speak for both Gary C. and myself on that.

As ever even more, Greg