9203A

Date: Sun Mar 01, 1992 3:19 pm PST
Subject: Re: Auroral Activity Warning

From [Avery Andrews (9203011519)]

I don't think they call them anything, since they're hardly ever seen (we are much further from the South Pole than you are from the North)

Avery.Andrews@anu.edu.au

==
Date: Mon Mar 02, 1992 8:31 am PST
Subject: Stir it up

[From Rick Marken (920301)]

Well it's been pretty quiet out there. I hope that means everyone is working hard. I have been. I've finished maybe 5 pages of "the book" -- not much but at least I was concentrating on it. I didn't catch the aurora -- but I saw one when I was up there in the North -- good ol' Minnesota. But one incredible astronomical phenomenon per year (the July eclipse for me) is enough for me.

I hope we get more of the Beer/Powers stuff soon. I don't know about others, but in only two or so posts so far the difference between the PCT approach and the conventional approach is becoming crystal clear -- and in a very detailed way. Note how the Beer approach relies on stimuli to GUIDE behavior. In the PCT approach, we try to imagine what sensory variable the animal might be controlling if we are to see a particular behavior. The PCT approach would be a lot easier if there were data on the kinds of variables that bugs actually control. Now we are in the position of doing "the test" in a kind of roundabout way; we guess at the controlled variable based on our understanding of what the organism "does" (what we see as its actions) and what we know of how the environment works (the latter being the toughest part). Then we make the model control these variables in this kind of environment. This approach is problematic because lack of success could be a result of 1) poor choice of controlled variable or 2) insuffficently detailed environmental model. The best way to do this modelling, I think, is to start by determining some controlled variables.

One last infuriating comment -- related to jury duty. I have not yet been selected for a jury -- but the two I was almost selected for were cases where no crime was committed -- well, a law was broken but there was no crime from my point of view; nothing was hurt other than some sensibilities. Do we really need to spend all this time and effort trying to stop people from violating our good taste?

By the way, one was a prostitution case; the other a drug case. There are people begging on the streets, shooting each other regularly, robbing, hurting, abusing and being abused. And

I'm supposed to waste my time deciding whether or not someone
put the wrong kind of thing into their own mouth. Jeez.

Ok, this isn't PCT but, then, we had so much fun
with religion, maybe we can get some excitment stirred up
on victimless crime.

Best regards Rick

==

Date: Mon Mar 02, 1992 9:44 am PST
Subject: Crime and PCT

[from Gary Cziko 920302.1030]

As Rick Marken awaits his fate as a possible jurist, his recent comments (920301a) have got me thinking about crime and punishment.

I have never been called for jury duty. I don't know why, since I do vote. The last time (about a dozen years ago) I thought I might be called I figured there was no way the prosecution would ever select me because I was a determinist in the Skinnerian sense. People did what they did because of the contingencies of reinforcement provided by the environment. How could I call someone guilty and justify punishing him or her when I believed that I or anybody else on the jury (or the judge for that matter) would have done exactly the same thing if given the same genome and raised in the same environment? We were all beyond freedom and responsiblity.

Now, PCT has given me new way to look at behavior, but I haven't thought much of what PCT means for dealing with crime and justice. What does it mean for someone to be "guilty?" Does it mean having a "bad" system concept and/or principles and following them? Or does it mean having a "good" system concept and/or principles but NOT following them? Neither perspective seems satisfactory to me.

Rick, have you been thinking about these issues (for "real" crimes, not for victimless ones)? I think that this is a different question from the kinds that Hugh Gibbons has used PCT for. Once you do get selected, perhaps you can give us some insight on what the jury is controlling for (I can see you as a natural-born disturbance in such a situation).

--Gary

===

Date: Mon Mar 02, 1992 1:58 pm PST
Subject: Beer on Powers Bug

[from Randy Beer 920302 via Gary Cziko to CSGnet]

Dear Bill,

I apologize for my delay in responding to your last message but, as I mentioned in my last message, it is difficult to find the time to actively participate in discussions such as these.

As near as I can tell from your description, it sounds like your proposed feeding system should work. One clarification. At several points in your message, you state that one advantage of your proposed control system over the current appetitive and consummatory circuits is that it would allow the insect to move as food is eaten or moved out from under it. The implication seems to be that the present circuits do not have this capability. But, in fact, they do. Indeed, in my original simulation, food patches shrunk as they were consumed, so this capability is crucial. If the insect loses contact with the food while it is eating, the consummatory controller immediately relinquishes control to the the appetitive controller, which will reorient the insect to the patch. This can be seen in Figure 8.7 of my book. If the insect loses contact with the food, the consummatory command neuron will stop firing and its inhibition of the search command neuron will disappear until contact is reestablished.

I seem to detect a strong aversion to having "superordinate" control systems turning off subordinate control systems in your comments. Perhaps I would understand this aversion better if I was familiar with HCT. However, I should point out that, to take just one example, the consummatory controller doesn't literally turn off the appetitive controller when it is activated, it simply denies it access to the motor system for turning. The relative odor strength on each side of the insect is continuously computed. However, this odor strength can only affect the insect's turning when the search command neuron is enabled.

In fact, I don't really think of the appetitive CIRCUIT as being subordinate to the consummatory CIRCUIT at all, since portions of both are always active. Rather, the appetitive BEHAVIOR is subordinate to the consummatory BEHAVIOR. I tend to think of the artificial insect's nervous system as consisting of a number of partially overlapping circuits "fighting" for control of the periphery. Which circuit wins at any particular time depends upon both the external stimuli and the internal state of the insect (consider, for example, the interaction between edge-following and the appetitive phase of feeding). There is some (rather tenuous) neuroethological evidence for this way of looking at things, but as you probably know, work on the neural basis of behavioral choice is still at a rather primitive stage.

More generally, I'm afraid that, due to my inexperience with this group, I don't have a very clear idea of the purpose of this discussion, though I sometimes have the vague sense that you are arguing for or against something. Are you interested in how nervous systems control behavior? Are you interested in designing artificial autonomous agents? Are you interested in how the artificial insect works? Are you interested in a rational reconstruction of the artificial insect using HCT? Each of these topics lead to a somewhat different perspective on the issues we have been discussing. If I better understood your motivation, then perhaps I could tailor my comments accordingly.

Best regards,
Randy Beer
==
P.S. from Gary Cziko

Randy Beer is not on CSGnet. Responses to him should therefore be sent to his personal address <beer@cthulhu.ces.cwru.edu> with a copy to CSGnet.

My own preference right now would be for Beer's involvement with CSGnet to be limited to interactions with Powers, lest we lose Beer completely as the result of heavy salvos from CSGnetters.

But who am I to tell anybody on CSGnet what to do (other than the fact that as "listowner" I can cut anybody off from the net whenever I wish and there just ain't nothing nobody can do about it, so there!)?--Gary

Date: Mon Mar 02, 1992 4:09 pm PST
Subject: BEERBUG reply

[From Bill Powers (920302.1500)]

Randy Beer (920202) --

Don't feel in a rush to answer anything. I have plenty to do, too, and quite understand.

>... you state that one advantage of your proposed
>control system over the current appetitive and consummatory circuits
>is that it would allow the insect to move as food is eaten or moved
>out from under it. The implication seems to be that the present
>circuits do not have this capability. But, in fact, they do.

I agree, they do.

>I seem to detect a strong aversion to having "superordinate" control
>systems turning off subordinate control systems in your comments.

It's a mild aversion. I like to see as much as possible emerge from a model without being explicitly put into it. I don't think you would like a model in which the direction from the animal to the food acted to turn the animal in that direction (i.e., if the food were at azimuth 60 degrees, an ad-hoc neuron would rotate the animal's body by 60 degrees). I think the underlying idea is to try to keep the modeler's intelligence out of the model as completely as possible (an ideal which, I realize, all modelers are forced to violate frequently for lack of data or clever enough designs). The model should operate strictly from its own organization, using only information available to it in the form of sensory signals. I think your model does fairly well in this regard, although some parts of it (particularly obstacle avoidance and edge following) seem a bit awkward to me (one of my scientific criteria).

With regard to the "search" neuron, I just think the model would be prettier if it didn't have discrete modes. Maybe they're unavoidable. But by using the approach circuitry I suggested, you can leave the food-approach system turned on all the time, with hunger bringing it automatically into action, and proximity to food stopping the forward motion without any need to recognize the abstract condition "Now I'm at the food and don't need to move any more." In the same way, "wandering" doesn't have to be a special condition any more: when the animal is satiated (energy error is zero), other motivations can make it move, with the noise level in the food-direction sensors providing the random bends in the path.

By leaving all these systems on all the time, you can get some interesting (and I think real) conflict situations, and see how the bug resolves them. In my model of crowd behavior, the individuals use goal-seeking and obstacle-avoiding circuits that run concurrently, and the resulting behavior as a "person" finds a way through a crowd of other people to a goal looks far more intelligent than it actually is. The "people" will even retrace paths to get out of traps. That's the effect I like to see: minimum design, maximum apparent complexity of behavior.

So you can see that my scientific criteria for models rest in good part on such subjective notions as niceness, prettiness, and simplicity. This applies, of course, only when the data don't constrain us to a particular design.

I'm still working on simplifying the locomotive system. The data on the real cockroach's gait, p. 82, show that the forward swing wave begins as the leg moving backward on the other side passes the midpoint of its travel during the stance phase (this seems to be true at all speeds). A circuit that detects this midpoint position of the rear leg and generates a pulse can trigger the swing wave on the contralateral side. Can I assume a sensor that responds continuously to leg angle, or do I have to do this open-loop using the motor driving signal? I'm going to make leg angle proportional to the driving signal, by the way, rather than using a pseudo-force output with some rather odd physics involved. If leg angle sensors exist (other than the limit sensors), I can make leg angle a controlled variable.

If the triggering of a swing wave can be made automatic, then the speed control circuit becomes very simple: just a time integrator. I don't know if this is going to work out, but it looks promising so far.

My swing-wave generator will produce ALL gaits, with the tripod gate as a natural limiting case at high speeds. It continues to work at all speeds down to zero. You were right, by the way, in pointing out that more than one leg can move at a time on one side: I didn't look closely enough at the diagram.

Reversal of direction is going to be interesting. In real cockroaches, is the swing phase still initiated in the rear legs while traveling backward? Or does it start at the front?

>Are you interested in how nervous systems control behavior? Are you
>interested in designing artificial autonomous agents? Are you interested
>in how the artificial insect works? Are you interested in a rational
>reconstruction of the artificial insect using HCT?

All of the above, but the emphasis is on modeling the behavior of real, particularly human, organisms. The "hierarchical" aspect of the modeling may be less important in simple organisms, because the goals are going to be pretty simple and there won't be many levels of organization. More important is the concept of control -- the idea that the system VARIES its actions to CONTROL variables defined by its input apparatus. As an example, instead of thinking of the food patch as causing turning behavior via the odor sensors, think of varying the direction of motion as controlling the unbalance of the sensor signals in an odor gradient field. This shifts the viewpoint from the observer-centered laboratory system to the organism-centered control system. In simple systems this shift of viewpoint doesn't make much difference, but when things get complex it can make the behavior much more understandable. It also helps one avoid putting too much of the observer into the model -- when you think of control of perception, the world that matters is the one represented by sensor signals, not the one the observer sees. Most of the "behavior" we observe is a more or less irrelevant side effect of what the control system is really doing.

When I get a working model, by the way, I'm just going to dump it in your lap to do with as you please. I have no desire to publish in this field. I hope you'll see some principles in it that will interest you. I'd like to see you start using the CT orientation in your work, but I have no urge to compete in your area. I figure the best way to recruit you is to demonstrate the CT approach using something dear to your heart like your pet cockroach. When we get past this preliminary stage, I think you'll start seeing some real power in the CT approach. Then you can start teaching us things about real neurons in real control systems.

Best regards, Bill Powers

Date: Mon Mar 02, 1992 6:05 pm PST
Subject: Crime and PCT

[From Rick Marken (920302)]

Well, just my luck. A mistrial was declared so I don't get to be on a jury once again. But I'm on lunch break so I'll make some comments on Gary's post.

Gary Cziko (920302.1030) says:

> I was
>a determinist in the Skinnerian sense. People did what they did because of
>the contingencies of reinforcement provided by the environment. How could
>I call someone guilty and justify punishing him or her when I believed that
>I or anybody else on the jury (or the judge for that matter) would have
>done exactly the same thing if given the same genome and raised in the same
>environment? We were all beyond freedom and responsiblity.

Yes. I don't know why Skinner didn't call for the immediate dismantling of the entire legal system in the country since, if he were right, the courts would be, at best, irrelevant.

>Now, PCT has given me new way to look at behavior, but I haven't thought
>much of what PCT means for dealing with crime and justice. What does it
>mean for someone to be "guilty?" Does it mean having a "bad" system
>concept and/or principles and following them? Or does it mean having a
>"good" system concept and/or principles but NOT following them? Neither
>perspective seems satisfactory to me.

A person is guilty if they intentionally produce results that are perceived equivalent to the results descibed in words (laws) as being forbidden. I have no problem with the law here. It is illegal for me to make my car go faster than 55mph on the freeway. If I produce that result (even accidentally) I am guilty. Some results must be produced intentionally for there to be guilt -- so the law must do a verbal and circumstantial version of "the test" to establish intentionality. This is an interesting study in itself and maybe

PCT could make some contributions to the law in terms of determining the intentionality of an act in retrospect.

I think the deeper question is "what are laws" and why do we have them? This will overlap a lot of the old social control discussion. But right off the top of my head I think laws are just an attempt to agree on where people should set their reference levels for certain results. Now there's going to be all kinds of detailed problems with this -- since people will not necessarily perceive the world in the same way, or describe it with the same words, let alone agree on a reasonable reference level for anything. But, basically, laws are just an articulation by a group of where their common references should be and where these references probably already are set if everyone is generally living fairly successfully in the group. Thus, the biggies (like murder, theft,etc) are pretty much already agreed on -- but, as the ten commandments shows, people seem to feel better when it is publicly stated that "this is our reference level for X".

I think laws are a good thing -- they make people feel like they are agreeing about variables of common concern. But the making of laws is an on going process because, as societies change there are changed results that can occur, new ways to produce results, etc. So there have to be changes in our references all the time -- and groups of people must decide, all the time, what are the "best" referece levels to agree to. I don't know how "best" is determined but when it is determined then we get new laws.

Now the problem, as I see it, is that there seem to be many laws that do nothing but create crime. I don't know why these stay on the books -- but they are generally what I call victimless crimes. My favorite example is drug laws. I can understand that most people don't want to perceive drug crazed people roaming their streets. So a law that says "no drugs" seems sensible enough. Except that many people find that setting their own reference for drugs to the level required by law is difficult -- it must produce inner conflict. So they get the drugs anyway. Any this is a process that creates REAL problems; violence, theft, etc etc. So I don't understand
why people don't see that the law (even though they like the reference it represents) is a bad thing. They did during prohibition. My conclusion:laws are not designed to solve problems!!! (at least not problems like daily murder and robbery that is a direct result of having a law against drug posession).
They are designed, as I said, to reflect consensual reference levels -- and if this consensus results in massive societal crisis then tough. Catastrophic consequences of consensual references don't seem to make much difference to people -- and this makes sense from a control theory perspective. People just want to have things match their refernce -- and I think people have a reference for having laws that mirror their own references.
Most people have a reference for "zero drug use" . The law reflects that and so most everybody is happy -- even though our cities are battlefields. This is not a rare occurrence.
The Catholic church still maintains as it's own law that birth

control is forbidden -- even though the side effects of this are theoretically and observably catastrophic. But its more important to have a law that matches that "respect for life" reference of whatever it is.

So I guess I am proposing that people would apparently want,as laws, things that reflect their reference for the way particular perceptions should be (zero drugs, zero prostitution, zero [insert the thing that repulses you the most here]) even if putting the coersive power of the state behind that reference level means catastrophe for the society. I believe that most of the laws that have this effect (nice idea but bad side effects) are crimes where the only victim is the perpetrator. So it gets me a bit upset. If drug laws, prostitution laws, etc really made things better -- and had none of the horredous side effects that their enforcement produces -- I'd support them as strongly as I support the freway speed limit law (even though I got a ticket one year ago).

I think these "victimless crime" laws are just another reflection of people's inability to keep from trying to control other people. I think control theory is obviously relevant here.

Best regards Rick

Date: Tue Mar 03, 1992 8:45 am PST
Subject: conference in Denver

Date: Mon, 2 Mar 92 17:39:57 EST
From: Davi Geiger <geiger@medusa.siemens.com>
To: <many addresses deleted>

CALL FOR PAPERS

NEURAL INFORMATION PROCESSING SYSTEMS (NIPS)
-Natural and Synthetic-
Monday, November 30 - Thursday, December 3, 1992
Denver, Colorado

This is the sixth meeting of an inter-disciplinary conference which brings together neuroscientists, engineers, computer scientists, cognitive scientists, physicists, and mathematicians interested in all aspects of neural processing and computation. A day of tutorial presentations (Nov 30) will precede the regular session and two days of focused workshops will follow at a nearby ski area (Dec 4-5). Major categories and examples of subcategories for paper submissions are the following;

Neuroscience: Studies and Analyses of Neurobiological Systems, Inhibition in cortical circuits, Signals and noise in neural computation, Theoretical Neurobiology and Neurophysics.

Theory: Computational Learning Theory, Complexity Theory, Dynamical Systems, Statistical Mechanics, Probability and Statistics, Approximation Theory.

Implementation and Simulation: VLSI, Optical, Software Simulators, Implementation Languages, Parallel Processor Design and Benchmarks.

Algorithms and Architectures: Learning Algorithms, Constructive and Pruning Algorithms, Localized Basis Functions, Tree Structured Networks, Performance Comparisons, Recurrent Networks, Combinatorial Optimization, Genetic Algorithms.

Cognitive Science & AI: Natural Language, Human Learning and Memory, Perception and Psychophysics, Symbolic Reasoning.

Visual Processing: Stereopsis, Visual Motion, Recognition, Image Coding and Classification.

Speech and Signal Processing: Speech Recognition, Coding, and Synthesis, Text-to-Speech, Adaptive Equalization, Nonlinear Noise Removal.

Control, Navigation, and Planning: Navigation and Planning, Learning Internal Models of the World, Trajectory Planning, Robotic Motor Control, Process Control.

Applications: Medical Diagnosis or Data Analysis, Financial and Economic Analysis, Timeseries Prediction, Protein Structure Prediction, Music Processing, Expert Systems.

The technical program will contain plenary, contributed oral and poster presentations with no parallel sessions. All presented papers will be due (January 13, 1993) after the conference in camera-ready format and will be published by Morgan Kaufmann. Submission Procedures: Original research contributions are solicited, and will be carefully refereed. Authors must submit six copies of both a 1000-word (or less) summary and six copies of a separate single-page 50-100 word abstract clearly stating their results postmarked by May 22, 1992 (express mail is not necessary). Accepted abstracts will be published in the conference program. Summaries are for program committee use only. At the bottom of each abstract page and on the first summary page indicate preference for oral or poster presentation and specify one of the above nine broad categories and, if appropriate, sub-categories (For example: Poster, Applications-Expert Systems; Oral, Implementation-Analog VLSI). Include addresses of all authors at the front of the summary and the abstract and indicate to which author correspondence should be addressed. Submissions will not be considered that lack category information, separate abstract sheets, the required six copies, author addresses, or are late.

Mail Submissions To:

Jack Cowan
NIPS*92 Submissions

University of Chicago
Dept. of Mathematics
5734 So. University Ave.
Chicago IL 60637

Mail For Registration Material To:

NIPS*92 Registration
SIEMENS Research Center
755 College Road East
Princeton, NJ, 08540

All submitting authors will be sent registration material automatically. Program committee decisions will be sent to the correspondence author only.

NIPS*92 Organizing Committee: General Chair, Stephen J. Hanson, Siemens Research & Princeton University; Program Chair, Jack Cowan, University of Chicago; Publications Chair, Lee Giles, NEC; Publicity Chair, Davi Geiger, Siemens Research; Treasurer, Bob Allen, Bellcore; Local Arrangements, Chuck Anderson, Colorado State University; Program Co-Chairs: Andy Barto, U. Mass.; Jim Burr, Stanford U.; David Haussler, UCSC ; Alan Lapedes, Los Alamos; Bruce McNaughton, U. Arizona; Barlett Mel, JPL; Mike Mozer, U. Colorado; John Pearson, SRI; Terry Sejnowski, Salk Institute; David Touretzky, CMU; Alex Waibel, CMU; Halbert White, UCSD; Alan Yuille, Harvard U.; Tutorial Chair: Stephen Hanson, Workshop Chair: Gerry Tesauro, IBM Domestic Liasons: IEEE Liaison, Terrence Fine, Cornell; Government & Corporate Liaison, Lee Giles, NEC; Overseas Liasons: Mitsuo Kawato, ATR; Marwan Jabri, University of Sydney; Benny Lautrup, Niels Bohr Institute; John Bridle, RSRE; Andreas Meier, Simon Bolivar U.

DEADLINE FOR SUMMARIES & ABSTRACTS IS MAY 22, 1992 (POSTMARKED)
please post

Date: Tue Mar 03, 1992 3:14 pm PST
Subject: PCT bug -- spread it around

From Pat & Greg Williams (920303)

>Bill Powers (920302.1500)

>When I get a working model, by the way, I'm just going to dump it in your
>lap to do with as you please. I have no desire to publish in this field.

We hope you'll dump it in our lap, too, so we can include some genuine PCT circuitry with Version 4 of NSCK. (You'll get -- the usual meager -- royalties, of course....)

Best,

Pat & Greg

Date: Tue Mar 03, 1992 5:22 pm PST
Subject: speed limits

I sent the following privately to Rick Marken, thinking it not an issue for CSG-L, but he urged me to send it to the list. So here it is.
=================

You say:
> If drug laws, prostitution laws,
>etc really made things better -- and had none of the horrendous
>side effects that their enforcement produces -- I'd support
>them as strongly as I support the freway speed limit law
>(even though I got a ticket one year ago).

I think that the freeway speed laws could be added to your list of laws that set references that people like, but that have catastrophic consequences (well, very mildly catastrophic, compared to the drug laws). Over the last decade, I probably have driven more in Europe, largely Germany, than in N. America. In Germany, the lack of speed limit on most stretches of the autobahn leads people to drive carefully and with due regard to what happens around them. On some stretches, there is a speed limit, which might be anywhere from 100 kph to 130 kph. On these stretches, traffic clogs (which might be a cause rather than an effect of the speed limit), people often drive in lanes other than the right lane even when they are not passing, people (occasionally) pass on the right when a righteous speed-limited driver hogs the left lane, and generally you get a highway mess like those to which we are accustomed.

In France or Italy, the highways are officially speed-limited, but in practice they are not (speed traps occur, but are rare on freeways). Driver behaviour is good, as a rule, though maniacs exist (as here).

In all my driving, I never feel as scared as in the first days after I return from driving comfortably at 160 kph in moderately dense traffic to our local highways where most people keep below 130 (speed limit 100), but have no lane discipline or consideration for each other.

So--I agree whole-heartedly about the victimless crime laws, but I think there are many laws relating to acts that might have victims, that also have results opposite to the effects their proponents intend. Speed limits on freeways is one such, I think. Capital punishment is another.

Private mail--this isn't a CSG issue.

Martin

The Jury's still out, I suppose.

Date: Tue Mar 03, 1992 7:42 pm PST
Subject: Subject: Black Box - RKC

from Bob Clark

to Bill Powers.

Re: the 4-port black box, your note of February 11. Of course one can describe the situation in terms of its being "affected by the state of some external variable." However, if one is curious about some particular box, it must be intentionally "disturbed" if one is to discover the nature of one (or more) external variables it may be controlling. While discussion in terms of "states of variables" is certainly possible, I find a more experimental approach more useful. I also think some people are unfamiliar with the word "state" in this sense.

You suggest a difference between the case where there is an external link between terminals #4 and #2 versus an independent variable applied to #2. The experimenter, being outside the box, can observe any external connections. And any disturbance from the experimenter is (really, by definition) an "independent variable."

I am intrigued by the question of the boundary between the Box and its environment. To the experimenter, it is quite clear. The environment includes everything except what is inside the Box. But what does the Box perceive? Its Behavior is the Control of its Perception. But what does it perceive? Only the disturbances of those connections that cross the walls of the Box. And its output actions reflect the differences between disturbances affecting terminals 2) and 3). In addition, some kind of connection must exist outside the Box, not perceptible by the Box but observable by the experimenter. If these relations are not found, the Box does not act as a Negative Feedback Control System.

These considerations seem to me to be applicable to many non-living systems, and the use of non-living systesm as extensions of human systems is impressive.

Enough for now. Regards, Bob Clark

Date: Wed Mar 04, 1992 8:05 am PST
Subject: fuzzy logic, perceiving language

[From: Bruce Nevin (Wed 92014 10:35:52)]

An article on fuzzy logic in _New Scientist_ for February 8, pp 36-9, has intrigued me. It's been several years since I read about Zadeh's work. The perspective of HPCT has helped me to understand it in a new way.

People are using neural nets to define fuzzy sets. These set definitions--which resemble your "subjective probabilities," Martin--are then provided to implementations of fuzzy-logic inference engines. In a simplified example, a wash load has degrees of membership in fuzzy sets involved in two (out of 12) inference rules, as follows:

```
     Wash Load    *     Cloth Quality   ==>        Washing time
     =========          =============              ============
     0.52 "Average"      0.39 "Soft"               0.39 "Short"
     0.28 "Heavy"        0.70 "Generally Soft"     0.28 "Average"
                                                   ------------
                "Center of Gravity"                4.9 minutes
```

A "defuzzification" procedure that computes the "center of gravity" as a definite number is mentioned rather than described in the article.

It seems to me that HPCT maintains what might be called virtual fuzzy sets in real time rather than by this batch-mode external production of them in a neural net system. Putting it this way provides access to a currently very hot button.

To point this up, the box on p. 38 of the article describes Takeshi Yamakawa's fuzzy controller and its solution of the "inverted pendulum control problem." A shaft is mounted on a vehicle by a pivot. The vehicle has angle and velocity sensors for the shaft and velocity sensors for the vehicle. The problem is to keep the shaft vertical by moving the vehicle, much as one would balance a stick on the palm of one's hand. The HPCT solution is obvious. The fuzzy controller cannot be robust under disturbances such as tilting the table, turning on a nearby fan, etc. But this is thought of as real hot stuff, gets megabucks of support in Japan, and now elsewhere throughout the world.

The popularity of fuzzy logic stems probably from the craving to see the whole world from the Principle level, in terms of programs, no?

Can we show how HPCT maintains the virtual equivalents of fuzzy sets and does the equivalent of fuzzy inference for this sort of control? Conversely, are fuzzy categories not needed on the program level? Consider here Bill's (920225.2230) example "I've been rich, and I've been poor, and believe me rich is better." Is this not just what you were talking about, Bill? The equivalent of fuzzy sets and fuzzy logic instead of discrete categories and binary either/or logic.

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

Chuck Tucker (920225) --

>citations of studies that involve disturbances and disruptions within
>interaction of two or more persons

Brad Goodman did some work here 8-10 years ago involving an instructor and a learner separated by a screen, much as Bill has discussed. I have looked for some old BBN reports describing this work but haven't found it yet. I'll keep looking, but meanwhile here is a possible lead for you to follow up.

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

Avery (2/14/92, 2/15/92) --

Your explanation to David about why you think the Diverian approach is all post hoc rationalization" did not explain. Do diverians reject in principle any attempt to predict "what sentences are grammatically acceptable and what they mean" in a way that could be modelled on a computer? Or do they reject rewrite rules manipulating phrase classes (which I also do)? Or something in between?

Generative grammar of course does not provide any coherent account of meanings, nor does LFG as a particular flavor of generative grammatical theory. So the goal stated above must be reduced to predicting "what sentences are grammatically acceptable" and leaving out differences of acceptability on semantic grounds. Are you then saying that Diverians reject the notion that grammatical acceptability is a central issue? In that they would indeed be closer to Bill's position, as I see it.

Avery (27 February 1992) --

Looking at your example:

```
                S
             /     \
        /            VP
    /              /    \
 NP:SUBJ         /          \
  NOM          V             NP:OBJ
   N         broke            NOM
 Jimmy                         N
                           the glass
```

Notice that you get an identical structure for

(1) Jimmy's betrayal broke up their marriage.

```
                S
             /     \
        /            VP
    /              /    \
 NP:SUBJ         /          \
  NOM          V             NP:OBJ
   N         broke up         NOM
 Jimmy's                       N
 betrayal                  their marriage
```

To distinguish these, you have another S subordinate to NP for (1), and special metarules that say that rules like S ==> NP VP apply except when the S is so subordinated, when other rules apply instead. Or in TG you have rules operating on trees of phrase-class labels resulting from the first sort of rules, producing deformed trees under these conditions. In operator grammar, the reductions apply directly to words. All those abstract names of phrase classes are unneeded and only get in the way of connecting the words with the perceptions.

Corresponding to (1) we can also say:

(2) Jimmy's act of betraying someone broke up their marriage.
(3) Jimmy's betrayal broke up their state of being married to each other.
(4) An act broke up a state; the act was Jimmy's betrayal of someone; the state was of their being married to each other.

These all exist. Regardless of the likelihood of one or another being said, they all are English sentences conveying the same meanings (quibbles possible about word choices, e.g. "state" vs. some other word, but that's what Webster's advocates for "marriage" so I went with it.) They are not related to each other in any obvious way in LFG or in GB theory or in any other theory using a phrase-structure-based system of rules. They are transparently related to each other in operator grammar.

Furthermore, the transitions (by minimal sentence-differences) from less explicit and less regular sentences to others that are more explicit and more regular constitute the pathway by which one may arrive at a semantic representation for sentences. After all the work of LFG, or GB, or any of the other PSG-bound theories is done, a semantic representation still must be devised, and a mechanism relating semantic representations to words and syntactic structures on the one hand, and to perceptions on the other. Typically, the semantic representation includes features like [+abstract] to make the difference between N=glass and N=marriage. This additional vocabulary and the additional syntax governing its combinability has to be correlated with perceptions (meanings) at the same time as the ordinary words do. In operator grammar there are only words to be correlated with meanings. Those words and the dependencies among them (whether obscured by reductions or made explicit by undoing reductions, sometimes artificially) constitute the semantic representation that is to be correlated with meanings.

This pellucid simplicity is forever beyond the reach of theories that base themselves in PSG rewrite rules that manipulated the names of classes of phrases. Saying that N-bar and N-double-bar are really just instances of N, for purposes of determining the head of a construction, is formalistic ad-hoc-ery. They are still different names of different classes of phrases.

Conservation of resources on a small serial computer is not a germane criterion if we are modelling human control of perception including language.

Bill Powers (920224.0800) --

>What's missing . . . is the sense of "Yeah, that's how I do it."

It appears that we keep the complexity of language control out of awareness for good reason. There is just too much going on at once. In order to describe it to ourselves and talk about it to one another we must use that which we would talk about. But using it and making it wiggle real slow under controlled conditions so we can get a good look are mutually exclusive.

>It could be claimed that the mental machinery
>is invisible to the observer and that all we get out of it is the result.
>But it seems strange to me that these methods use words, symbols, and rules
>for manipulating them in the familiar way, except for the fact that they
>propose content that is strange to me -- Bruce's expansions, and Avery's
>NP, VP, trees, recursions, etc.

I argue that the "expansions" of operator grammar do not change the content, only the form. The content is all familiar words, and the pairwise word dependencies are all familiar. Some of the larger string configurations are unconventional and unfamiliar. You get many unconventional but possible sentences with other theories. With other theories there is also new content, namely, a metalanguage vocabulary of NP, VP, etc. and its syntax of trees, recursions, etc.

To the extent that the unconventional sentences are sayable and understandable, however unlikely they may be in practice because of their awkwardness and their violation of convention (which is why they undergo reductions)--to the extent that they exist, they must be accounted for by any theory. Operator grammar exploits them to provide an explicit and informationally complete semantic representation for all sentences.

>I've tried for quite some time now (longer with Bruce) to elicit a
>description of what a linguist is doing in the process of getting from a
>received sentence to the structural analysis. What I get back is further
>analysis -- i.e., you DO it, but you don't DESCRIBE WHAT YOU'RE DOING. I
>don't want to know that "bite" is a word that takes two arguments, or that
>"bite" is a PRED function of SUBJ,OBJ. That doesn't tell me what you're
>doing in your head to get from sentences to those statements.

How a linguist gets from a received or imagined sentence to a structural analysis is called linguistics and it may possibly be relevant to how a language user gets from a received sentence to meanings. You say you want a description of the former when in fact you want a description of the latter.

>I don't want
>to know what a system or theory says about how the words are related --
>what I'm trying to get is a description of what happens in your
>consciousness when you begin with a new sentence (Josh felt Jean was
>unsympathetic) and begin to re-represent it or something about it. Instead
>of arguments developed from further application of the method in question,
>I'm trying to hear the processes going on in present time in the person
>offering those arguments.

You're asking for a freeze-frame account of category recognition, mostly. I don't think I can break into the black box that takes sensory input on one side and outputs a category on the other. I can introspect on the sensory input, but I can only speculate on the process by which it is categorized.

I look at the sentence (truncated from your last quoted above):

(5) I'm trying to hear the processes going on in present time.

I recognize I as a word that can't be an operator--it can't be said "about" some other word. This is a zero-order word, in OG parlance.

I recognize 'm trying as a reduction of am trying.

Ignoring here what am..-ing is a reduction of, I recognize try as a word that can be said "about" two other words. Two other words must be present (said or zeroed in a reconstructable way). These are its arguments. The first one has to be a zero-order word, preferably one to which I am attributing human characteristics. The first one cannot be an operator word. The second one cannot be a zero-order word, it has to be an operator word.

I recognize that I is probably the first argument word that must be present in order for try to be said. I recognize an operator-argument dependency between try and I. (Please allow me to use meta-words like operator and argument so as to talk about these things without imputing to me the claim that I use these words internally in carrying out the process I am describing.)

Furthermore, with try I have a strong preference that two zero-order words be the same, namely, the first argument of try (which is I) and the first argument of the operator word that is the second argument of try. I haven't identified the second argument of try yet, but this strong preference sets up some limitation as to what sort of word it might be: it must be something that could have I as its first argument.

Also under try, there is a strong preference for reducing the repeated first argument and the operator to the preposition to plus the operator.

I recognize the next word as the preposition to. Given the preceding, I recognize this as a reduction of I as argument of a following operator word. (This dependency could be made explicit by "expanding" to "I'm trying that I should hear" but we need not do that expansion to get at the dependency.)

I recognize hear as a word that cannot be said without two other words present. The first argument must be a zero-order word. The second argument must be an operator word. (When I say I hear John, I mean I hear John doing something. The process of and justification for reducing dually-classified words to a single class is a separate issue that we can take up if it troubles you.)

Given the preceding, I recognize that the second (reduced) occurrence of I is the first argument of hear.

I don't have time to continue this, but perhaps this gives the flavor.

Is this an accurate representation of what goes on in a language user in real time? I couldn't prove that. But I think a model could be built to do this in a convincing way. One's theory of language filters what it is possible to perceive about one's control of perceptions resulting in language. Having a least theory, that imports the least structural baggage of its own, is I think of critical importance. Alternative theories with which I am familiar are over-structured.

I've spent too long on this. Got to run.

Bruce
bn@bbn.com

Date: Wed Mar 04, 1992 11:41 am PST
Subject: fuzzy logic, Gatherings

[From Rick Marken (920304)]

Bruce Nevin (Wed 92014 10:35:52) says:

>To point this up, the box on p. 38 of the article describes Takeshi
>Yamakawa's fuzzy controller and its solution of the "inverted pendulum
>control problem."

>Can we show how HPCT maintains the virtual equivalents of fuzzy sets and
>does the equivalent of fuzzy inference for this sort of control?

I think the best comment on this fuzzy control stuff was made about a year ago when the fuzzy "inverted pendulum controller" was posted to csgnet. This program was huge (in source) -- many megabytes. Bill Powers (in a matter of hours) posted an inverted pendulum controller based on good old fashioned control theory. It was maybe twenty lines of code, most of which handled the physics. My opinion -- fuzzy logic is a solution looking for a problem. The problems that it has been applied to already have better solutions. But fuzzy logic is trendy science -- so that's where the bucks go.

I think you are right -- people seem to want to deal with the world at the principle level or higher. I see this here at Aerospace where there is an unbelievable reverence for "expert systems" -- and this is coming from engineers, many of whom have been trained in control theory. I think they are so used to dealing with control problems verbally -- like "That voltage there is decreasing TOO QUICKLY. We better switch to SIDE B"-- that they don't notice that they are just controlling quantitative variables (keep the rate of dischage = x). It's a very interesting phenomenon, this regression to symbol manipulation by quantitatively trained people. I think it speaks to how deeply imbued we all are with stimulus-response preconceptions about behavior.

To Bill Powers: You may have done this before but I would appreciate it if you could post a simple diagram of the "persons" in the Gatherings program. I want to write a program that demonstrates (on an individual basis) some of the points you make with the gatherings program -- i.e. behaviors like "circles of people" and "backtracking" are just observed side effects of each person's efforts to control peceptions like proximity, speed, etc. So what are the controlled variables in the Gatherings program and how are they related? I think I remember that one was "proximity" and the other is "speed" -- is that right? Are they controlled independently? Thanks for the info.

Regards   Rick

Date: Wed Mar 04, 1992 2:18 pm PST
Subject: Psych Review Article

[From Rick Marken (920304b)]

I forgot to thank Mark Olson for calling attention to the Psych Review article by Vallacher and Wegner and for his comments on it (it's nice to have you back Mark). I havn't gotten a hold of it yet but I plan to get it this weekend and look forward to some good laughs (er -- interesting reading).

> They suggest that
>somehow their model is better[than PCT], but I don't see how or why.

Well, maybe they are right. Boy, would THAT be a surprise.

Hasta Luego   Rick

Date: Wed Mar 04, 1992 3:08 pm PST
Subject: Re: fuzzy logic, perceiving language

A couple of years ago we discussed fuzzy control here. I'm a researcher in possibility theory and generalized information theory. Bill and I both came to the conclusion that "fuzzy control" is effectively "fuzzy S-R control", another form of non-linear standard control. Yes, it appears to be hot stuff for people.

Actually, fuzzy theory extends far beyond the relatively simple control theory applications it's been mostly used for. Furthermore, the idea that fuzzy sets ARE linguistic categories is VERY limitting, but instilled by Zadeh himself, and very hard to shake.

> People are using neural nets to define fuzzy sets. These set
> definitions--which resemble your "subjective probabilities," Martin--are
> then provided to implementations of fuzzy-logic inference engines.

Fuzzy researchers are always pointing out that fuzzy membership grades are NOT probabilities in any way, shape or form.

> Can we show how HPCT maintains the virtual equivalents of fuzzy sets and
> does the equivalent of fuzzy inference for this sort of control?

To my understanding, no. My understanding of HPCT is that percpetion and action are both crisp things. Input is a number, RL is a number, error is the difference. These have no UNCERTAINTY. Is there a PROBABILISTIC version of HPCT? Can perceptions be WEIGHTED so that e.g. temperatre is perceived as "kind of high, kind of low"? If so, then there can be a FUZZY HPCT.

> Consider here Bill's (920225.2230) example "I've been rich, and I've
> been poor, and believe me rich is better." Is this not just what you
> were talking about, Bill? The equivalent of fuzzy sets and fuzzy logic
> instead of discrete categories and binary either/or logic.

Contrary to what some fuzzy people would have us believe, not all linguistic or psychological cateogories need be represented by fuzzy sets, and vice versa. Above you're contrasting Rich and Poor as two DISCRETE and DISTINCT categories: one cannot be BOTH rich and poor. Uncertainty is about relaxing the excluded middle, about allowing shades of grey. We could have weights: .4 rich, .6 poor. Here .4 + .6 = 1, so these could be probabilities. If we had .8 rich and .4 poor, then we have fuzzy membership grades, or POSSIBILITIES.

Really, fuzziness and control theory have little to do with each other. Fuzziness is another way to represent uncertainty, whether in a control system or elsewhere. Control theory is about stability despite disturbances, whether uncertainty is involved or not.

O-->
| Cliff Joslyn, Cybernetician at Large, 327 Spring St #2 Portland ME 04102 USA
| Systems Science, SUNY Binghamton   NASA Goddard Space Flight Center
| cjoslyn@bingvaxu.cc.binghamton.edu   joslyn@kong.gsfc.nasa.gov
V All the world is biscuit shaped. . .

Date: Wed Mar 04, 1992 3:12 pm PST
Subject: Re: Psych Review Article

[Martin Taylor 920304 17:00]
(Rick Marken 920304b + Mark Olson)
>
>
>I forgot to thank Mark Olson for calling attention to the Psych Review
>article by Vallacher and Wegner and for his comments on it (it's nice to
>have you back Mark). I havn't gotten a hold of it yet but I plan to get it
>this weekend and look forward to some good laughs (er -- interesting
>reading).
>
>> They suggest that
>>somehow their model is better[than PCT], but I don't see how or why.
>
>Well, maybe they are right. Boy, would THAT be a surprise.
>
I see no such claim in their paper. The only mention of Powers is in a list
of people who have proposed hierarchic theories, unless I missed something.
What they do claim is moderately interesting and in no way incompatible with
PCT, though not supportive either. It may bear on the old discussion of the
function and place of consciousness.

A gross simplification of their claim is that if the actions that allow high-level control are easy, then what subjects see themselves as doing is what we would call satisfying the high-level reference. But if the lower level control structure is disturbed or not well structured (the actions are more difficult), then people see themselves as "doing" the low-level things. It's the difference between "visiting Aunt Dorothy" and "trying to find the doorbell (for Aunt Dorothy's door)" (Not their example). Their claim also is that the behaviours that the subjects identify themselves as doing are less susceptible to disturbance than are the other levels. If someone is "keeping fit" then rain may cause them to switch from jogging to indoor exercise, whereas if they are "jogging" then rain is less likely to stop them. (Like, but not the same as, one of their examples).

In PCT terms, the idea seems to translate as:
(1) If an ECS is maintaining good control, then it is not high in conscious awareness, and (2) the conscious identification of a reference signal serves to stabilize that reference signal. The reference signal itself seems to become a percept for another "conscious" control system.

In a sense, the reference must be a percept, because the person holding it describes themself as perceiving it to be one. But it is also functioning as a reference signal at the same time, in the normal way. It seems to me that this idea alters the standard wiring diagram in an interesting way that might relate to the issue of conflict and reorganization.

Thanks for pointing out the reference, Mark.

Martin

Date: Wed Mar 04, 1992 3:19 pm PST

Subject: More on zeros

[Martin Taylor 920304 17:30]

I'm still have a problem with the zeros in a perfectly controlling hierarchy. Rick has sent me a SYLK version of his spreadsheet that I have not yet downloaded, so I will hold off more comment until I have tried to analyze better its behaviour. But Bill triggered something a couple of days ago with his comment that to stably control something like walking around a circle involved a lot of changing references that resulted in a lot of actions.

What this triggered was a thought about what it means to control continuously for a sequence. In simple terms, suppose one is controlling for a sequence of values between 0 and 1, and the percept strays from the desired values in a consistent way. Then an error signal is presumably generated, which affects lower-level control systems in such a way as to reduce that error, but it does so only for future parts of the sequence. If they stay close to the desired values, the error signal should likewise stay close to zero. But the sequence, when it is finished, was not the one desired. Shouldn't that result in an error signal? Or, are there no error signals during the acquisition of the sequence, and only a *multi-degree-of-freedom* error signal at the end, which could affect subsequent attempts to perceive the desired sequence? Neither view seems satisfactory.

Looking at it from the more abstract view, the sequence can be described by a z-transform of the time series. The behaviour of the (sampled) controller can be described in the same terms. The bheaviour of the desired percept (variable according to factors outside the ECS) is entirely merged with the intrinsic characteristics of the controller-environment loop. Most such cases are dealt with by methods such as inverse filtering, turning the time-spread signal into a series of point events, but I cannot see that this approach would be useful here. It would be akin, I think, to the second possibility listed above, to report only a whole-sequence error after the fact.

In respect of the zeros problem, sequence provides a distinct glitch, in that the reference sequence is not a simple representation of a one degree-of-freedom (scalar) error signal. It provides a set of arbitrarily changing references based on some internal memory to the lower ECSs. These then are subject to the transients that I excluded in my original questions about the hierarchy in close control of slowly changin disturbances.

Rick's spreadsheet doesn't have sequences, if I remember correctly, and if it nevertheless shows non-zero percepts and references throughout the hierarchy after the initial transient has decayed, then I have to rethink my intuitions (if that concept makes sense). Until then, I don't want to proceed with the degrees of freedom part of the discussion.

I may not have time to do it anyway. I have three chapters for different books to complete before April 1, only two of them so far drafted and none completed. Then I will be away for about 2 months (as yet indeterminate) and out of touch with CSG-L. So we may have to postpone the debate until summer.

Martin

Date: Wed Mar 04, 1992 4:59 pm PST
Subject: Fuzzy Logic; language

[From Bill Powers (920304.1000)]

Bruce Nevin (920304.1035) --

I hope all unusual circumstances are being coped with; nice to hear from you again.

Fuzzy Logic:

I saw a demo of the fuzzy logic inverted pendulum controller. Pretty poor control, as you mention. If there's a place for this approach, it has to be at the higher levels -- I certainly wouldn't use it in place of analog variables and quantitative continuous controllers (in other words, you're right about its applicability to my description of "richness"). I have a suspicion that fuzzy logic represents a step away from the purely digital concept of brain function, caused by the discovery that logical systems don't control continuous variables very well. Unfortunately, analog concepts of simulation aren't taught along with digital concepts: if you're adept at digital logic, you probably "don't do analog." If students became adept at both approaches, I think they would find it much easier to know when analog is appropriate, and when logic is needed for simulating higher-level processes.

Language:

>It appears that we keep the complexity of language control out of
>awareness for good reason. There is just too much going on at once.
>In order to describe it to ourselves and talk about it to one another >we
must use that which we would talk about. But using it and making it
>wiggle real slow under controlled conditions so we can get a good look
>are mutually exclusive.

This is what I was asking about. The problem here is that there isn't any experimental or observational test to see whether the proposed invisible processes are actually occurring unless you specifically, and slowly, make them occur. You seem to be saying that the expansions are really taking place, in words, but REAL FAST and not in awareness, during ordinary speech. How do you tell whether this is a right, as opposed to merely plausible, description of what is going on?

I think we can boil this down to a specific dispute. What I DO experience is that a terse sentence results in incomplete perceptual meanings and that I then complete them to eliminate dual or contradictory meanings. Having done that, having perceived the ambiguity or incompleteness and having remedied it in imagination, I then have many choices as to how to describe the more complete meaning in detail.

You appear to be saying that the terse sentence is first recognized as incomplete in terms of linguistic rules -- this word requires two arguments and only one was supplied, and so on. Then the phrase is expanded to supply the arguments, after which the phrase evokes the completed or unambiguous perceptual meaning. An alternative is that linguistic analysis results in recognition that a word is missing, after which a search for words with acceptable meanings is done, selecting a missing word appropriate to the meanings already evoked. In either case, the linguistic analysis occurs before the completed meaning is experienced.

When I hear a sentence like "Rino Sanders hit safely," I don't think that I first realize that "Rino Sanders hit" requires another word. I think that the first result is an image of Ryne Sanders hitting -- swinging a bat -- and (when the object of hitting isn't mentioned, or even if it is mentioned, before the name of the object is heard), the crack of the bat against A BALL, the scamper for first base, the fielder getting the ball back too late ("hit safely"), and so on. Bingo, the whole little scenario pops into place and I know what the words mean. I don't detect any intervening linguistic analysis of any flavor. In fact, if the sentence is "Rino Sanders hit a car," I have to retract the whole immediate perceptual interpretation and supply a different one, with a sensation of a big error at the instant the word "car" is heard.

You seem to be saying that between "Rino Sanders hit .." and the appearance of a completed perceptual meaning in my head, there is an ultra-rapid invisible linguistic process that supplies a missing word in order to satisfy the requirement of two arguments, and only then does the image of a BALL enter my experience.

I would argue that the missing ball is supplied first, and that the motivation for supplying it is that the meaning of hitting is incomplete without the image of something doing the hitting and something being hit. I think you are proposing that the missing word is noticed first, the motivation for supplying it coming from the general requirement that the operator in question is known to require two arguments, only one of which has been supplied.

Is this what we're arguing about?

Sentence deconstruction:

>>(5) I'm trying to hear the processes going on in present time.

>I recognize I as a word that can't be an operator--it can't be said
>"about" some other word. This is a zero-order word, in OG parlance.

Does this recognition come before the meaning of "I" as the sense of a person?

>I recognize 'm trying as a reduction of am trying.

Does this recognition come between hearing "I'm trying" and getting sense of a person trying to do something?

> ... I recognize try as a word that can be said "about" two other >words. Two other words must be present (said or zeroed in a >reconstructable way).

Since the "other word" here is a whole phrase, "to hear the processes going on in present time," or at least "to hear," are you saying that you get no meaning out of "I'm trying" until something being tried is imagined or heard?. Try it on this: "I'm wondering .... what you mean." During the dotted pause, doesn't "I'm wondering" convey any perceptual meaning?

>I recognize that I is probably the first argument word that must be
>present in order for try to be said. I recognize an operator-argument
>dependency between try and I.

What would you think of this:

I immediately suggests a person; try immediately suggests the person in process of doing something unspecified. This knowledge, that something is not yet specified, is immediate and perceptual. In a formal linguistic analysis, we say that the person and the trying are to be classified as an operator and an argument. The sense that something is still missing, perceptually, suggests to us that another argument is called for by this particular operator-word. This in turn suggests a rule: some operators require two arguments. The logic then follows: this is such an operator; it has less than two arguments; therefore there is a missing argument to be supplied.

Here's an analogy without the linguistic analysis: You see a picture of a man leaning far forward holding onto a rope that goes backward over his shoulder and then straight horizontally out of the frame. Without talking about it, you know that there is something out of the frame resisting the pull of the rope. Maybe the other end is tied to a post, or a donkey, or another man pulling the other way. So you can see the man, and you can see that he is pulling, but you can't see what he is pulling on. Nevertheless, you know that he is pulling on SOMETHING, via the rope. If the man is progressing forward, you wait in anticipation of seeing what the rope will pull into the picture.

Now you could create a description of what you see: the man is pulling with the rope. The omission of WHAT he is pulling is clear, not because of the words, but because of the picture and its wordless sense of something on the other end of the rope resisting the pull, which you imagine. You could say "The man is pulling ...." and wait for whatever is on the other end of the rope to appear before you name it: "... an elephant!" Or you could use a word that refers to a missing perception, an error signal: something. What led you to expect something to appear? Was it the knowledge that the name of this action is an operator requiring the name of an agent and the name of an object? Or was it the wordless comprehension that there was something outside the picture resisting the pull? I suggest it was the latter: you can't know what kind of operator "pulling" is until you see its perceptual meaning. There's nothing in the word itself to indicate how many arguments it takes, or even whether it's an operator or an argument.

Before you can say whether an operator-word requires one argument or more, you have to know what specific perception is meant by it. "Shout", if it refers just to a vocal spasm, takes no argument. "The man shouted hello" is clearly an operator that takes two arguments. It isn't the word, but the perceptual meaning that "takes arguments." And the arguments "taken" aren't words, but perceptions of configurations, events, and so on that are related by the action described. Words are classed as operators and arguments because of the way their meanings are related; their meanings are not related as they are because the words are classified as they are, but because the perceptual world works the way it does.

This leads me to an observation that's been a-building for some time while I lacked a way to talk about it. It's the difference between a descriptive rule and a prescriptive rule.

Here's a neutral example of a descriptive rule: Bodes's law. Take the numbers 0,3,6,12,24 ... (doubling each time after the three). Add 4 to get

4,7,10,16,28 ... . Divide by 10; the numbers are very nearly the distances of the planets from the Sun in astronomical units (Earth's distance = 1.00). When Uranus was discovered, it was found at 19.19 A.U., where the rule of Bode's Law said it should be at 19.6 A. U. This led to a hunt for planets at the missing position, 2.8, between Mars and Jupiter, and the discover of Ceres, the first known asteroid. Pluto was discovered at 39.5 A.U., fitting Bode's Law which predicted it at 38.8. Unfortunately, Neptune comes between Uranus and Pluto (most of the time), between two numbers in the sequence.

This rule is very impressive, except that there's no known reason WHY the planets should approximate this series -- except perhaps that their distances are bound to approximate SOME series, which someone was bound to discover.

An example of a prescriptive rule: a knight moves two squares ahead or back and one sideways, or two squares sideways and one ahead or back. Why does this rule hold true? Because behavior in chess is deliberately adjusted to keep the moves of knights in conformity to this rule. If someone moves a knight in any other way, the opponent or any onlooker will cry "Illegal move!" and force the move to be corrected. We therefore observe that knights always move according to this rule.

Neptune does not "obey" Bode's Law, but this "illegal position" does not give rise to any action that makes it legal again. So Bode's law only describes. The movements of chess pieces do, literally, "obey" rules, for if a move violates a rule it will be set right. The rules of chess are prescriptive.

Now the question I have for Bruce Nevin and Avery Andrews is this: are the rules of grammar that you are helping to develop descriptive, or prescriptive? In short, is there an underlying system that forces language to exhibit these rules, or are these rules like Bode's Law -- interesting but fortuitous approximations with no necessary basis in the natural world?

It seems to me that if any set of rules is prescriptive, we will find that deviations from the rules are corrected. Would this not be a way of choosing between rival grammars?

Best,   Bill P.

Date: Wed Mar 04, 1992 8:46 pm PST
Subject: language

Re: Bruce Nevin (4 Mar 1992)

>Your explanation to David about why you think the Diverian approach is
>all post hoc rationalization" did not explain. Do diverians reject in
>principle any attempt to predict "what sentences are grammatically
>acceptable and what they mean" in a way that could be modelled on a
>computer?

My reading of Diverians is that they reject *all* attempts to systematically predicate the grammatical and semantic properties of individual sentences from antecedently specified principles.

What they do instead is `validate' analyses by somehow coming to the conclusions that the distributions of constructions in texts ought to be systematically skewed in various ways, and then doing counting & statistics to see if they are. I see this method as a useful supplement to other ones, but not much more than than. And it somehow doesn't seem PCT-ish at all!

On `Jimmy's betrayal':  why do you think I or any other Chomskyan needs a [+abstract] feature here?  All I think I need is (a) a semantic specification for `betrayal' (mostly shared with the verb `betray') (b) the info that the `treacher' argument can be expressed as a possessive (this in fact probably being predictable from more general principles). I'd say that `betrayal' is a noun and `betraying+NP' is a verb (with gerund inflection) on the basis of things like:

```
  the betrayal/ing of Mary cost John his happiness
. *the betraying Mary cost John his happiness.

  John's frequent betrayals of Mary were the scandal of the department.
  John's frequent betraying of Mary was ...
 *John's frequent betraying Mary was ...

 *John's betrayal of Mary frequently is shocking.
 *Johns betrayals of Mary frequently ...
  John's betraying Mary frequently ...
```

Chomsky has some more discussion of this sort of thing in `Remarks on Nominalizations'.

>Generative grammar of course does not provide any coherent account of
>meanings, nor does LFG as a particular flavor of generative grammatical
>theory.

I think that Jackendoff provides an account which is perfectly coherent within its own terms, though it leaves out things that I and my philosophical friends think ought to be there (e.g. the objective and social aspects of meaning). & I have this suspicion that the fact that we continually interact with the world is also important for the theory of meaning. But I don't see how different flavors of grammatical theory could have much to do with this.

Bill Powers (920304.1000):

>Now the question I have for Bruce Nevin and Avery Andrews is this: are the
>rules of grammar that you are helping to develop descriptive, or
>prescriptive? In short, is there an underlying system that forces language
>to exhibit these rules, or are these rules like Bode's Law -- interesting
>but fortuitous approximations with no necessary basis in the natural world?

A bit of both is my answer, but it's a complicated issue that I'm actually writing a paper about at the moment. They are (a) descriptive because we don't know how the mechanisms that produce them work (b) more interesting than Bode's law because they can be seen to interact in complicated ways, so it is probably the case that they correspond in some systematic way to actual facts about brain structure (the Peacocke/Davies concept of `description at level 1.5', in case any of you are into the philosophical literature.

Bill Powers (920227.0900):

< on language development & development of control in PCT >

This requires a fair amount of actual work to answer, so the answer is I just don't know. There has been heaps of work (about which I know virtually nothing) on language & cognitive development, but presumably not using PCT categories (I'd guess that Piaget would be the closest one could get). And presumably people didn't notice the right things, so it would all have to be done all over again. My inclination is to ignore development for the moment & continue thinking about adult grammar & how to get a beer.

Avery.Andrews@anu.edu.au

Date: Thu Mar 05, 1992 4:24 am PST
Subject: a metaview of beer discussions

To: CSGnet people
From: David Goldstein
Subject: conflict
Date: 03/05/92

The recent "discussions" about Beer have made me review what PCT has to say about intrapersonal versus interpersonal conflict. What is the difference theoretically speaking?

As we have talked about before on CSGnet, intrapersonal conflict (within the same person) exists when two control systems are trying to put the same lower level perception in two different states at the same time. The PCT solution is to use the method of levels to raise awareness to a level above the control systems. The solution to an intrapersonal conflict is to view both control systems at the same time from a higher level and to reset the reference signal for at least one of them. A patient of mine likened this to pulling both ends of a rope.

As a concrete example, a worker may be furious at his boss for some specific reason in a specific situation. Self-image 1 of the worker may want to be masculine while self-image 2 may want to be socially adjusted. To control each of these self-images, different actions would be required. A superordinate "observer" inhibits(sets reference signal to zero) self-image 2 from being operant in the situation. The person acts to control self-image 1 which results in "socially adjusted" behavior.

Interpersonal conflict(between two people) seems like a different situation. There are some new options open for ending the conflict. One can escape the other control system(ignore, leave). (It is true that in dissociative disorders, for example, multiple personality disorder, a person switches to a different self-image and escapes the stress in this manner.) One can physicaly eliminate the other control system(fight, war). There is no higher level control system to reset the reference signal of one of the conflicting parties(except by law or other binding agreements).

In the Beer discussions, the conflict ended when one of the parties pointed out that the friendship between them was being endangered by continuing the discussions along the lines that it was going. Both parties stopped the discussions (almost). Trying to prove that each person's viewpoint was correct was less important than trying to remain friends. The conflict was ended (for now) by agreeing to control for a different experience, namely, being friends which is important to both parties.

When I think about cases of marital conflict which I have been involved in as a therapist, it seems to me that those couples who want to stay married are the ones who find alternative ways of handling differences and dissatisfactions.

It seems that the discussions are now heading in this direction. Each of the parties are showing new understandings of the other's viewpoint. If they completely adopted the viewpoint of the other, then the conflict between each of them would become an intrapersonal one. Maybe this is the way of solving interpersonal conflict. Make the other person's viewpoint your own. Then it gets solved as an intrapersonal conflict. If each party in the conflict comes up with the same or similar higher level maneuver then we have conflict resolution.

In short, the way to solve an interpersonal conflict is for each person to turn it into an intrapersonal conflict and then go from there. In the March 9, 1992 issue of Time, there was a description of Cyrus Vance and how he solves interpersonal conflicts between nations. His strategy is: "Master the facts of the situation; listen exhaustively to both sides; understand their positions; make sure they understand the principles that must dictate a solution; and don't give up."

Perhaps interpersonal conflicts require a third person because it is so hard to really understand the other person's viewpoint as if it were your own. The third person helps each person do this and also provides the higher level principles which any solution must satisfy.

I don't mean to give the impression that intrapersonal conflicts are that easy to solve. In the difficult cases, it may require the help of a second person, a therapist, who helps the person discover the higher level principles which are behind the conflict as well as to help the person identify and describe the nature of the conflict.

Date: Thu Mar 05, 1992 6:28 am PST
Subject: language

[From: Bruce Nevin (Thu 92015 08:22:03)]

Bill Powers (920304.1000) --

There's too much going on to track and report it all at once. I think we are perhaps each reporting what goes on in one hemisphere. Both aspects are necessary and each supports the other. And questions of which comes first are I think ill conceived. The two processes jostle and interfere with and support one another in pandaemonium parallel.

Another factor is sublanguage, specializations of language for particular subject-matter domains. A game such as baseball or chess, or a particular technical domain such as immunology or pharmacology (studies of which I have mentioned) or HPCT or sailing or auto mechanics, each such domain is characterized in part by a specialized form of the language used by its participants. These sublanguages differ both in syntax and semantics. Even the constitution of words may differ. I have mentioned that "the beating of the heart" is a single "word" of the Symptom class in a sublanguage of pharmacology, even though it is a phrase made up of many words in other usage, especially in the sublanguage of physiology whence it is borrowed into that of pharmacology.

So, "hit" has associated with it particular syntactic and semantic possibilities in a sublanguage of baseball that differ from those in other domains. If you hear "Rino Sanders hit the umpire" you must shift from the sublanguage of baseball to a kind of language appropriate for talking about a fight. Having done so, you don't expect to see Sanders placidly walk toward first base upon hearing "the umpire told him to take a walk."

Yes, this has very much to do with the world of input and imagined perceptions (I must never forget that), but it also has very much to do with institutionalizations of the perceptual world that depend in great measure upon learned linguistic patterning (you must never forget that).

I will try to respond more fully next week.

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

An initial paraphrase for clarity:

A descriptive rule is not enforced by anything. A violation merely shows that there are exceptions to the description, or that the description needs refinement. Your example is Bode's Law. Rules of this sort are regularities discovered about or attributed to nature.

A prescriptive rule is enforced by something. A violation is resisted, presumably by some control system. Your example is the definition of how a knight can move in chess.

My version: a knight can move from corner to corner of a 2x3 rectangle. He gallops across country. These are two different descriptive statements of the prescriptive rule.

In talking about language, the prescriptive/descriptive dichotomy is entirely in the realm of descriptive statements of prescriptive rules (in your sense of the latter).

A prescriptive<L> grammar is a description of prescriptive regularities in language that are socially esteemed plus an injunction that you should obey them, or a description of prescriptive regularities in language that mark one as socially inferior plus the injunction to obey.

A descriptive<L> grammar is a description of the prescriptive regularities in language. It should refer to the (sensu strictu) descriptive regularities in language as understood background. (These are never mentioned in a [linguistic sense] prescriptive<L> grammar precisely because they cannot be violated, insofar as the description in adequate. There's a minor epistemological rathole there, but a small one of rapidly diminishing returns, so we'll ignore it.) A descriptive<L> grammar should also describe the different socially-marked variants of the language and the values of each for the personal and social image of its users. A descriptive<L> grammar should also describe the subject- matter sublanguage specializations of the language and their intersections, borrowings (e.g. for metaphor like "take a different tack"), and other relations. A descriptive<L> grammar should describe similiar relations of borrowing and sometimes deeper intersection with different, perhaps unrelated languages used by neighbors, immigrants, invaders, merchants, and so on. A descriptive<L> grammar traditionally describes a time-slice of the language in its process of never-stopping change. However, by describing all these things, it describes the sources and bases for ongoing change processes. Furthermore, it must describe constructions that are marginal to the language at a particular time, but were more well established in the past, or will become more established and normal in the future. The synchronic/diachronic dichotomy is therefore artificial.

Needless to say, there are no descriptive<L> grammars that meet all these criteria.

But an important point is that the descriptive<L> grammar describes what is available to that other hemisphere (metaphorically speaking, but perhaps literally as well) that is controlling and imagining nonverbal perceptions, and it describes the institutionalizations embodied in language that in-form those processes of the control of perception. The descriptive<L> grammar or a model of these aspects of the control of language does not describe that world of perception or those processes of control. Another description or model does that.

Each of these models requires the other.

>are the
>rules of grammar that you are helping to develop descriptive, or
>prescriptive? In short, is there an underlying system that forces language
>to exhibit these rules, or are these rules like Bode's Law -- interesting
>but fortuitous approximations with no necessary basis in the natural world?

Language has both aspects. The descriptive (your sense) aspects are linguistic universals referred to as background for a descriptive<L> grammar. An example is the way acoustic properties of the speech apparatus define regions at which articulatory deviations make little acoustic difference, and at which formants are easier to identify because they are clustered or merged. These regions are the familiar labial, aveolo-dental, velar, and back velar articulatory positions. A descriptive<L> grammar describes prescriptive (your sense) regularities reflecting norms or conventions that constitute a particular language at a particular time. Within the constraints defined by prescriptive (your sense) factors or strong linguistic universals, nothing forces language to exhibit particular forms or rules except the historical contingencies of evolving human cultural institutions.

The existence of zero-order words and of operators classified as to the argument requirement of their argument words appears to be universal and I suspect is a byproduct of hierarchical perceptual control. The existence of reductions and their general types and even certain generalizations about their form also appear to be universal. Some aspects of discourse structure appear to be universal, though a great deal of work needs to be done even to find out what the issues are. Details of word shape and particular reductions of word shapes are historically contingent. Particular sublanguages, social valuations of language, and so on, are historically contingent and not universal, though some responses to these evidenced in historical change appear to be universal.

Mostly we don't know what the facts are, partly because people are still parroting the inspired guesses of brilliant pioneers as proven truths. For example, Roman Jakobson proposed that infants' babbling is a form of play that ranges over all the possible speech sounds of any possible human language, and that as they learn a language those specialized innate devices that are not needed atrophy. This turns out not to be true. Yet on this foundation was built an even more speculative edifice of innate language perception and production mechanisms for syntax and semantics as well as for phonology, such as that Pinker is talking about.

Avery --
Date: Thu, 5 Mar 1992 14:38:50 EST

I see your post in my mailbox now, have no time but will comment only very quickly now.

Thanks for the additional comments on Diver, no, statistical skewing does not seem PCTish at all.

>verb `betray') (b) the info that the `treacher' argument can

I don't understand "treacher".

I know about the kinds of examples that are in Remarks on Nominalizations, though it is more than 20 years since I first read it. Cross-paradigmatic communication is difficult, not least because both parties must be aware that this is what is going on and committed nonetheless to success in it. The second can't happen without the first, and I'm not sure you're yet aware that you have to buy on to the first if we are to talk productively here.

>from more general principles). I'd say that `betrayal' is a noun
>and `betraying+NP' is a verb (with gerund inflection) on the basis of

I didn't ask a question to which this is an answer. I asked how you distinguish betrayal and glass, both nouns. (I suggested a +abstract feature, as has been used in the past and is still used in some quarters.) Then I asked how you would describe the relationship of the betrayal-betraying-betray series of sentences. I said this relationship is transparent in OG and useful in getting to a semantic representation. I suggested that their relationship is not transparent in LFG or other PSG-bound theories, and that their relationship has no direct bearing on any sort of semantic representation in these theories. I asked (implicitly or explicitly, I don't remember which) whether these suggestions are true.

Gotta run.   Bruce

Date: Thu Mar 05, 1992 10:19 am PST
Subject: description/prescription

[From Bill Powers (920305.0900)]

Bruce Nevin (920305) and Avery Andrews (920305) --

There seem to be more ways of interpreting the "prescriptive-descriptive" question than I had anticipated. One of my meanings failed to get across, although all the replies were pertinent.

The missed meaning concerned whether alternate schemes for representing language at higher levels exist. They clearly do: Bruce recommends one, and Avery another. So these are alternative descriptions of orderliness that can be seen in language.

In one sense, these descriptive schemes can't both be right. That is, if one of them truly describes the way the brain does language, then the other is just a fortuitous ordering that continues to make sense but has nothing to do with the way the brain actually operates. As an outsider I don't care which one you accept as the "real" one and which as "fortuitous," or whether you decide they are both fortuitous. But at least one of them is fortuitous.

In another sense, both schemes can be right, although not both in the same person at the same time. That is, if a certain orderliness can be perceived in language, one can learn to create language in a way that prescriptively preserves that particular kind of orderliness, in addition to conveying desired meanings. One will simply not use constructions that violate the scheme or have no meaning within the scheme. So there would be certain constructions that Bruce would use but Avery would not, and vice versa. And of course there would be large numbers of constructions that both would use, each seeing that they fit his own scheme. Even error-correction could result in correcting errors under both schemes at once, although there would remain some errors that one person would correct while the other would not, and within those, some corrections that one would accept and the other would not: the Coin Game.

This problem, and I hope you both see it as a problem, becomes worse as the proponents of the different ordering schemes develop their own structures to apply to more and more instances of natural language use. The race is on to arrive first at the ultimate perfect ordering scheme that covers every known or possible sentence in every known or possible language. The unspoken assumption is that only one approach can succeed in doing this and that the other must fail. The corollary is that if one scheme does manage to bring all of language into a single orderly description, it must be the RIGHT description. But what if they both succeed, as I expect they ultimately will do? Doesn't this suggest that NEITHER of them is right in any objective sense?

And what of the poor non-linguist, who must produce a structured language without knowing EITHER scheme? Both proponents claim that they are describing processes that go on under the surface of language in every speaker. But two different processes are described. If both theorists can show that every natural language construction fits each theorist's scheme, what does this imply about the naive speaker? I think that the implication is, rather, about the schemes: it says that language is not objectively ordered in EITHER way, even though orderliness of each kind can be seen in language once you learn the rules of the scheme, and even though it is possible to order language production in conformity with either scheme.

I've been suggesting that the structure of language is derived in large part from the structure of nonverbal experience, where nonverbal experience is known, from the standpoint of the linguist, as meaning. PCT allows us to investigate the structure of nonverbal experience in the context of controlling it. To the extent that we can verify the controllability of entities at various levels in a discernible hierarchy of perceptions, we will know some structural constraints on the meanings that words are used to indicate, and these constraints must show up in language.

I don't deny that language conventions are superimposed on this basic structural influence and that these conventions can be studied in their own right. But I think that anything common to all languages will be found at the level of nonverbal experience, not in language conventions (except as these conventions are inherited from other languages). Every language, for example, will have a way of indicating agent, action, and object of action, because those are basic elements of perception common to all human beings: they are aspects of the experienced human world.

If we can trace certain structural constraints to the world of nonverbal perception, then they will no longer have to be explained in terms of rules relating words as words. This will render superfluous any aspects of rival schemes that are intended to derive these constraints strictly from linguistic considerations. When those aspects are removed, what is left of rival schemes may prove to be far less different than may seem now to be the case: in fact, the schemes may then be reconcilable.

If HPCT has anything to contribute to the discipline of linguistics, this is the kind of thing it will have to say.

Best, Bill P.

Date: Thu Mar 05, 1992 1:54 pm PST
Subject: Re: Psych Review Article

[From Rick Marken (920305)]

Martin Taylor (920304 17:00) commenting on the Vallacher and Wegner Psych Review article says:

>A gross simplification of their claim is that if the actions that allow high-
>level control are easy, then what subjects see themselves as doing is what
>we would call satisfying the high-level reference. But if the lower level
>control structure is disturbed or not well structured (the actions are more
>difficult), then people see themselves as "doing" the low-level things.

This does sound very interesting and relevent to control theory. I think your comments about consciousness of means vs ends depending on level of disturbance to be quite on target. Now I am even more interested in getting a hold of the article.

Regards   Rick

Date: Thu Mar 05, 1992 2:10 pm PST
Subject: The Ed Ford Show

Ed

I couldn't find your personal address for e-mail so I'm posting this on the net, but I don't mind if they all hear.

I got the copy of the tape. Thanks. I just had time for the beginning yesterday but I can hardly wait to see the whole show. It looks great and you are terrific. Great job.

Rick

Date: Thu Mar 05, 1992 2:51 pm PST
Subject: Re: The Ed Ford Show

Ed (direct):

Rick Marken said on CSGnet:

>I got the copy of the tape. Thanks. I just had time for the beginning
>yesterday but I can hardly wait to see the whole show. It looks great
>and you are terrific. Great job.

Now that we know it exists, I expect lots of people on CSGnet are going to want a copy. What do I need to do to get one?-   -Gary

Date: Thu Mar 05, 1992 9:40 pm PST
Subject: language

re Bruce (5 Mar 92)

`treacher' = betrayer

> I asked how you
> distinguish betrayal and glass, both nouns.

By giving them different semantic representations, with different argument-specifications (none for glass, various optional ones for `betrayal'). What I have in mind is hitching something like

Jackendoff's 1990 `conceptual structures' to an LFG, though I admit that I haven't actually implemented this yet.

The betrayal-betraying-betray sentences will then get similar semantic structures due to having similar semantics for their lexical entries (just like Jackendoff would do it). I get the impression that in the `standard' view (common to LFG, GPSG, and at least some GB work), lexical entries to a lot of the work that Harris would do by reductions.  In effect, rather than `expand' whole sentences as syntactic structures, one `expands' lexical items into their semantic structures, plugging arguments into appropriate positions (my LFG does do this, except that the `semantic structures' of lexical items are just strings with places to put indexes into, like the first argument of an fprint statement in C).

Avery.Andrews@anu.edu.au

Date: Fri Mar 06, 1992 7:11 am PST
Subject: Powers Reply
From Randy Beer

>I like to see as much as possible emerge from a model without being
>explicitly put into it.

Along these lines, you might be interested in some of our more recent work. We have been using genetic algorithms to evolve continuous-time recurrent neural networks for controlling the behavior of autonomous agents. For example, we have evolved a variety of locomotion controllers using a slightly more realistic variation of the artificial insect's body. (By the way, the artificial insect's physics isn't TOO odd. It corresponds to that for a highly damped system, which probably isn't too bad an approximation for legged systems)

>Can I assume a sensor that responds continuously to leg angle ...

Insect legs do possess sensors that respond continuously to leg angle and angular velocity (namely the so-called "chordotonal organs"), in addition to the more limit-like sensors that I incorporated into my model (inspired by "hairplate receptors" found near the leg joints in the cockroach).

>In real cockroaches, is the swing phase still initiated in the rear legs
>while traveling backward? Or does it start at the front?

Very little work has been published on backward walking in cockroaches. I believe that both of the scenarios that you mention can occur.

>I'd like to see you start using the CT orientation in your work [...]
>I figure that the best way to recruit you is to demonstrate the CT
>approach using something dear to your heart like your pet cockroach.

Aha! So this is your real motivation. Now the character of some of your comments makes a little more sense. However, rather than haggling about the asthetics of various details of the artificial insect, why don't you lay out your theoretical principles so that we can discuss them. I realize that it may be difficult to summarize your position in a few paragraphs. To be fair, I will attempt to sketch my current theoretical position below:

To begin, I should clearly state that I am interested in BOTH understanding, through computer simulation, the neural basis of animal behavior and in designing versatile and robust autonomous agents. The artificial insect was primarily derived from the latter interest. It was an attempt to demonstrate that neurobiological control principles could be beneficially applied to artificial agents. Happily, I stayed close enough to the actual biology that certain aspects of the artificial insect have also generated a great deal of interest among neuroscientists. However, the specifics of that project do not exhaust my current theorectical position.

My current position stems from the observation that the only thing directly selected for in evolution is external behavior. Natural selection does not care at all what's inside an animal as long as the interaction of the complete package (nervous system, body and environment) is such that the animal survives to reproduce. This is why I am somewhat skeptical of any a priori committments to particular internal organizations. Of course, what is inside an animal is subject to a variety of biochemical, development and historical constraints.

In order to simplify the discussion, let us assume that the environment and body are given a priori. The problem then becomes one of finding the right internal dynamics so that, when this dynamics is coupled to the body and environment, behavior necessary to the survival of the agent is produced. There are two extreme cases for this dynamics which are worth pointing out. First, if the agent has NO sensory inputs from the environment, then its dynamics must essentially model the dynamics of its environment (nearly impossible for realistic environments). In my opinion, most classical AI systems, which base their action on manipulations of internal representations of their situation, are of this purely "model-based" character. At the other extreme, we have an agent with rich sensory inputs but no internal dynamics. Such an agent is essentially a functional map from inputs to outputs. Such agents are constantly "pushed around" by their environments. They cannot take any initiative in their interactions and therefore exhibit no true autonomy. In my opinion, the current "reactive" movement in AI (e.g. behavior-based robotics, situated activity, etc.) is mostly of this character.

I believe that the correct mix must lie somewhere in between. An agent can build into its internal control dynamics a fair amount of the dynamics of its body and environment. However, this must be done in a flexible enough way that the specifics of its immediate situation, as communicated by its rich sensory inputs, can appropriately bias its internal tendencies. Interestingly, many of the dynamical neural network controllers we've evolved using genetic algorithms seem to exhibit this mixed character, i.e. some of the dynamics necessary to achieve a given task performance is directly incorporated into the network itself and some only emerges in interaction with the environment.

Some of these ideas may sound a bit different from the artificial insect. However, I should remind you that that work is over three years old and theoretical positions evolve. I feel no particularly strong committment to most of the specifics of the neural circuits that I designed, which is why I feel somewhat uncomfortable defending them to you. Of course, from a scientific perspective, these circuits are interesting insofar as they accurately capture the relevant neurobiology. However, from an engineering perspective, they are simply means to an end, namely that of achieving the appropriate dynamics of interaction between the insect and its environment. I think that there are some good reasons for at least familiarizing ourselves with biological solutions as we attempt to design appropriate controllers, but the biological solution isn't necessarily the only one, or even the best one.

By the way, if you succeed in developing interesting elaborations of the artificial insect, I would encourage you to publish it. I have no fear of competition. Indeed, since there are far too few people working in this general area, I would strongly encourage you to do so. There is more than enough interesting work to go around.

Regards,
Randy Beer

Date: Fri Mar 06, 1992 10:48 am PST
Subject: Powers Reply

[From Rick Marken (920306)]

Hi Randy, welcome to CSGNet. I'm the loose canon on the west coast. I am trained as a cognitive psychologist and my main interest in control theory is as a model of human nature.

Let me take the liberty of giving a brief reply, from a psychological perspective, to the following proposal that you made to Bill Powers.

>   However, rather than haggling
>about the asthetics of various details of the artificial insect, why
>don't you lay out your theoretical principles so that we can discuss
>them. I realize that it may be difficult to summarize your position
>in a few paragraphs.

I think I can summarize the basics of PCT (perceptual control theory) in a few sentences. PCT begins with the observation of a phenomenon -- control. What we observe is that the events we call "behavior" are consistent results of variable actions produced by an organism. We also observe that the variabilty of actions is what is required to compensate for other variable influences (disturbances) that would eliminate the consistency of these results were it not for the nearly simultaneous effects of the actions of the organism. Thus, behaviors are controlled results of the combined effects of an an organism's actions and independent (environmental) disturbances. PCT is a model of how this control occurs. The central principle is that controlled results are perceptual representations of environmental variables that are maintained at internally specified reference levels (behavior is the control of perception). The actions that maintain perceptions at their reference levels are part of a closed, negative feedback loop which contains the appropriate dynamics (the theory defines "appropriate") so that controlled results are stabilized at their fixed or varying reference specifications. In the process of controlling perceptions an organism will produce many results besides those that are under control. These results may be interesting or important to an observer but they are irrelevant to the behaving system itself. Thus, the design of artifacts that aim to imitate the behavior of living systems must begin by determining the results of actions that the real system actually controls.

Simple as that.

Now let me ask for some clarification of your position:

You say:

>In order to simplify the discussion, let us assume that the
>environment and body are given a priori. The problem then becomes one
>of finding the right internal dynamics so that, when this dynamics is
>coupled to the body and environment, behavior necessary to the
>survival of the agent is produced.

What are "internal dynamics"? If, by this, you mean "a model of the nervous system" then I'd say that is certainly a way to describe the problem as I see it as well, though I would not take words like "behavior" for granted. Would the acceleration of the bug as it falls off a ledge count as a behavior to be modeled? If not, why not. If so, why so?

I hope you find time to continue to participate in CSGNet.

Regards

Rick

Date: Fri Mar 06, 1992 12:36 pm PST
Subject: Beer and CSGnet

[from Gary Cziko 920306.1400]

Rick Marken (920306) in response to Randy Beer's post said:

>Hi Randy, welcome to CSGNet.

Uh, not quite. Randy Beer is not on the CSGnet list which means he does not receive CSGnet mail. But since Bill Powers felt that the interaction he was starting with Randy would be of general interest to CSGnet, he encouraged Randy Beer to post his responses to CSGnet (anybody can post to CSGnet, whether officially on the network or not; but only people on the net get CSGnet mail).

This means that anyone wanting to respond to Randy Beer must send to his personal address <beer@CTHULHU.CES.CWRU.EDU>. By including both csg-l and beer as addressees, the message will go to both Beer and and CSGnet.

I will forward Rick's post to Randy this time but will expect others wanting to reach Randy to use his personal address. If Randy does decide to become an official CSGnetter, I will let you know.

Actually, this arrangement has certain advantages. It allows us to communicate to each other out of Randy's range and then communicate with him when we wish. But it means having to tack his address on anything meant for him.--Gary

Date: Fri Mar 06, 1992 1:10 pm PST
Subject: psych review article

Yes, this article is Very compatible with CT and doesn't claim specifically that the model presented is better than Powers' model, but in one paragraph near the end they make a short statement that their model is better than the "above" models, of which Powers' model is a part--that's why I said what I did.

Date: Fri Mar 06, 1992 1:54 pm PST
Subject: About My Tape

from Ed Ford (920306.14:20)

Gary - concerning the tape:

I was most fortunate to have been asked two years ago this month by the program manager for our local PBS station, to do a series of shows on relationship building through the television media. At that time, he was taking a course from me on marriage & family (loaded with lots of control theory) at a local renewal center. After submitting a proposal and having it accepted, product development got the project under way, the account executives got corporate sponsors, and I was asked to submit a script. Bill was kind enough to review my approach concerning control theory and offered valuable suggestions.

The last two weeks before filming had me doing "run throughs" almost daily and the show was filmed on Saturday, Jan. 11th, with 25 technicians, a make-up artist, a still photographer, three cameras, a beautifully designed set, cables everywhere, two technical TV trucks, an audience of 125 people, and me. Needless to say, I was somewhat overwhelmed.

It then went into "post production" where everything was professionally tied together, including graphics (including an explanation of CT), titling, audience shots, editing of unnecessary words, various audience shots and sounds strategically inserted, and now is ready for the premiere Monday, March 16th at 7 p.m. locally in Phoenix. The show was designed for the pledge drives, using my book Love Guaranteed as a promotion gift (which also explains control theory and includes in the back the reference list of CT books & papers that Bourbon and Bill Williams put together for my book, Freedom From Stress). If the ratings are good, other PBS stations will use it during their pledge drive weeks so it does have a chance of eventually being seen in other areas of the country.

The show is in two segments (24 & 21 minutes) and the first is partially dedicated to explaining CT. Also, the credits at the end mention the origin of CT as coming from William T.Powers, his books & the Control Systems Group. I tried to explain the concepts as simply as I could while still maintaining the essence of control theory. I'll have some in a few weeks for commercial sale for $20 plus $2 shipping. Paid up CSG members will be given a 50% discount, or $10 plus $2 shipping.

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

Date: Fri Mar 06, 1992 8:37 pm PST
Subject: BEERBUG; design & philosophy

[From Bill Powers (920306.1800)]

Randy Beer (920306) --

>We have been using genetic algorithms to evolve continuous-time
>recurrent neural networks for controlling the behavior of autonomous
>agents.

73 Ridge Place, CR 510 / Durango, CO 81301 -- in other words, I am interested. If I understand "genetic algorithm" correctly, the corresponding notion in PCT is "reorganization," which includes initial organization. That's considered an advanced subject because we don't know much about it. I'd like to see any papers you have describing the basic approach, or perhaps you wouldn't mind summarizing for CSGnet. If you don't mind, I'll leave that subject for another time -- there's plenty to say about it. This time I want to stick to basics.

> ... the artificial insect's physics
>isn't TOO odd. It corresponds to that for a highly damped system,
>which probably isn't too bad an approximation for legged systems)

Well, you have to admit it's a LITTLE odd, in that the whole bug is damped rather than the individual legs!

Design project progress report --

>Insect legs do possess sensors that respond continuously to leg angle
>and angular velocity (namely the so-called "chordotonal organs")

This is good news, because it means we can use analog control systems for leg position and leg velocity. Here's how it would work with semi-Beer neurons (internal details omitted):

```
 perceived leg vel        reference leg vel
        ---------->--------      |
       |                   |     |
       |                   X     O
       |                 -----------
       |                 \         /
```

```
       |                  \       /
       |                   \     /
       |                     \ /
       |     perceived        | (velocity error)
       |     leg pos          |
       |       ----->------   |                       O = positive input
       |      |            |  |                       X = negative input
       |      |            X  O position ref signal
       |      |          -----------
       |      |          \         /
       |      |           \       /
       |      |            \     /
       |      |              \ /
       |      |               |
       |      |               |   (position error)
    sensor    sensor       [MUSCLE]
  [VELOCITY] [POSITION]       |
       ^        ^             |
       |        |             |
       |        |             |
        -------[LEG]  -----<--
```

Actually this should be a balanced system, with pairs of oppositely-acting muscles, push-pull signals in place of single signals, etc. But the above will demonstrate the principles if you adjust constant current biases. The velocity control system should be biased to mid-range when the positive and negative inputs match. The lower system shouldn't need any bias in this simple circuit. Two simple Beer neurons should suffice.

The lower loop makes the sensed position match the reference position signal, which comes from the velocity error signal. I assume the leg is massless (some fast rate feedback would be needed for stability in the lower system with a leg having mass, if the internal muscle viscosity doesn't provide sufficient damping. I recommend leaving the leg massless for now!).

The lower system makes sensed leg position correspond to the reference input. The upper one makes sensed leg velocity correspond to the reference leg velocity, and acts by adjusting the reference signal for leg position. The gains of the two neurons should be high, so that only a small unbalance of the positive and negative inputs to each neuron is enough to produce a full-scale output change. The time-constant of the lower system should be very short.

If the reference velocity input is greater than midrange, the velocity of the leg will be positive; an input smaller than midrange will produce negative leg velocity. This gives us the basic circuit for moving the bug at a controlled velocity. This velocity will be almost independent of external forces: the bug will go at the same speed uphill, downhill, or on the level, either forward or backward. The same velocity reference signal enters all six leg control systems.

An auxiliary circuit, which I'm still working on, is needed to detect when the legs reach a forward or backward limit. Depending on whether the velocity is positive or negative, passing a limit will force the net position reference signal high or low, causing the leg to reset very quickly and overriding for a moment the velocity control system's output.

The reset signal will, in either case, also lift the leg. I'm planning to use the angle signal as one way of detecting limits, with the limit-sensing hair cells as a second way. Normally the limit sensors will trigger the reset, but if they fail, the angle signal will also trigger resets.

This is a completely symmetrical design: the leg will step forward or backward depending only on whether the velocity reference signal is higher or lower than midrange. The higher-level systems that use this 2-level system to control velocity therefore don't need to be concerned with resetting the legs or with doing any complex gating of connections to change from forward to backward motion.

When a reset is triggered, the effects will propagate (on each side of the body) to trigger a reset in the next leg rostrally, if any. Each rear leg's angle signal, when it crosses zero, will trigger the reset in the contralateral leg, thus making each pair of legs walk in the proper alternating pattern. The net result should be to reproduce the proper gaits at all speeds.

I think that the six sets of reset circuits will end up having a family resemblance to the arrangement of the six P-neurons in your present model. I'm not trying to use any special- or multiple-purpose neurons in this design; just the basic one. So many of the details in my circuit will probably prove to be collapsible into functions using fewer but more complex neurons.

You will note that this is an analog system, as shown. The reset circuits will use some digital logic, mainly set-reset flip-flops with some AND and OR logic to determine which way the reset should go -- all done with basic Beer neurons, of course.

The leg-lifting circuits can also be control systems. At the moment, I'm thinking that the appropriate controlled variable would be sensed pressure of the foot against the ground (or sensed load in some manner). I assume such sensors exist (?). To lift the leg, the reset system would make the reference pressure slightly negative. A very nice spinoff is that another higher system could then sense height of the body above the ground (a combination of leg-angle signals would do it, or perhaps hair sensors under the body, or even the visual system).

This higher system would compare sensed height with a height reference signal, with the output error signal adjusting the mean reference signal for all six foot-pressure control systems. As legs go up and down, the body height would tend to change, and the height control system would adjust the pressure reference signals upward and downward to maintain a constant height. Now the bug can walk over uneven terrain, for each leg will descend (after a reset) until the sensed pressure rises to the specified level -- and that will occur at whatever leg elevation puts it in contact with the ground. Furthermore, adjusting the height reference signal would make the bug walk with its body at an adjustable height off the ground. The bug could then creep under a door. This is pretty conjectural, though -- I haven't looked into the details at all.

---I hope you'll try the above analog circuit just to get a feel for how it works. Once nice thing about it is that you can use real physiology and real physics if you want to get into the complexities. In your model, leg velocity is determined completely by the speed of the body under a forward force, implying that the bug is walking in a viscious medium. In the revised model above, velocity is neurally controlled, using those angular-velocity sensors, and the muscle system and physics of body movement can be as realistic as you like (complete with a model of the muscle). There's no real need to go into such details, but with the analog model, you know you aren't violating any physical realities. Or neural ones, either, I hope.

So let's talk philosophy for a while.

>My current position stems from the observation that the only thing
>directly selected for in evolution is external behavior. Natural
>selection does not care at all what's inside an animal as long as the
>interaction of the complete package (nervous system, body and
>environment) is such that the animal survives to reproduce.

Agreed. But there is a long distance, and I believe several levels of organization, between natural selection and a walking cockroach.

Also, I think we have to be careful about what we mean by "behavior." In PCT, we distinguish between actions (outputs) and their consequences. Natural selection can't have selected for any specific pattern of leg movements in a cockroach. Those leg movements have to depend not on evolution but on where the food is. The consequence of moving the legs (under the right circumstances) is (from inside the cockroach) to bring food within eating-distance. So if evolution has selected for anything, it has to be for the consequences of moving the legs, not for those movements themselves.

To go even further, evolution hasn't selected for particular consequences, either, at the level of behavioral organization. What it selects for is a viable organism. Selection of _particular_ consequences is one of the things that the organism itself does; evolution has selected for organisms capable of doing that. This means that evolution has selected for control systems, because control systems control consequences of their actions, not their actions.

Look at the leg position control system above. The "action" of this system is to tense a muscle (or make one muscle more tense and an opposing one less tense). There are several consequences of this action. One is to move the body of the cockroach opposite to the leg movement, if the leg is down. The other, unconditionally, is to alter the signal coming out of the position sensor.

Note what is sensed: not the body "movement," but the position of the leg relative to the body. That is what is sensed, and that is what is controlled. This little control system makes the signal representing leg position match the reference signal it is receiving from a higher system. It is controlling not the muscle contraction, but a consequence of the muscle contraction.

This is easiest to see when disturbances occur. If some force is applied directly to the leg, the net force acting on the leg will change: the muscle still exerts the same force, but now we have added an external force. If this force pushes the leg in the same direction that the muscle force acts, the sensor signal will change, indicating too much movement. The negative feedback will result in reduction of the muscle force, and the leg will not, in fact, move (much). The greater the gain in the position-control neuron, the less effect the disturbance will have. If the disturbance aiding the muscle force is large enough, the muscle force will reverse!

When a disturbance comes along, the action changes equally and oppositely, so the action is certainly not under control by the bug. What remains the same, or nearly the same, is the sensor signal, the consequence of muscle activity AND the disturbance. This system makes one consequence of the action match the reference signal. Varying the reference signal will very reliably make the consequence -- the position -- change in the same way, even though the amount of muscle force produced can't be predicted without knowing every disturbance that is going to occur. It is the contention of PCT that evolution has selected for this kind of system.

In very simple organisms, it's possible that normal disturbances have so little effect on consequences of actions that no feedback control is required. If, for example, the cockroach's muscles produce large opposing forces, and if the stiffness of the muscle "springs" is very high, then contraction of the muscle will simply move the leg and the body. In that case, the lower system in the diagram could omit the position sensor and the negative feedback connection.

Likewise, if the organism were so small that the viscosity of the medium in which it moves completely determined the result of applying a force to it, then (as in your model) no velocity feedback would be needed either: the reference signal would just be a velocity command signal, and the output action would just be a force.

The presence of both velocity and position sensors suggests to me that feedback is needed, on this body scale. It is probably not needed by a bacterium, which coasts only about 1% of its own length when the flagellae stop spinning. But I suspect that cockroaches are big enough to need the feedback control. Of course this question will be settled when someone finds a way to do the right experiments.

You will notice that if the lower system works well, an external disturbance will have little effect on the second-level system. Only a very large force disturbance can make the leg be in a position other than the one that the higher system is specifying via its output, which becomes the lower-level reference position signal. For large enough disturbances, the second system will detect a velocity error, and adjust the lower level reference signal accordingly. This is the basic relationship of HPCT -- Hierarchical Perceptual Control Theory.

There are still higher-level control loops in your model. One of them, for example, controls the unbalance of signals from the odor sensors in the antennae. This unbalance is maintained at a reference level of zero. Your model doesn't provide for a variable reference signal, but it could. The same principle applies: the action of the system (now defined at a higher level of organization) changes when something disturbs the balance of signals from the receptors -- a disturbance could occur, for example, if the target food-patch were moved. So the action, the bending of the path, is not under control. What is under control, what becomes resistant to disturbance, is the unbalance of the odor signals, which is a consequence of the turning of the path and also of any changes in the location of the food.

Evolution couldn't select for the specific way the bug turns left and right. That turning has to be variable, in response to environmental changes. All that evolution could select for is the control system that is capable of keeping the signals from the two receptors balanced, a variable that has obvious bearing on the cockroach's ability to survive to reproduce.

I hope that all this gives you a clearer idea of what PCT and HCPT are about. The basis idea is that organisms control their inputs, not their outputs. The kind of organization that can do this is called a control system. We call it a _Perceptual_ control system to emphasize that only what is sensed can be controlled. We add the term "hierarchical" to refer to the way higher systems act by varying the reference inputs to lower systems. Systems of this kind control outcomes, consequences, not actions.

As to publishing, I don't need to publish for the sake of advancement or anything else I want, so I would be perfectly content if you were to take up these ideas and use them. If you don't become interested in them, maybe someone else will. My reference level is to perceive PCT being more and more widely adopted. My ego requirements will be met by a suitable mention in anything published. My hopes for advancing science would be satisfied if you were to begin thinking of yourself as a control theorist.

Best   Bill Powers

Date: Sat Mar 07, 1992 12:54 pm PST
Subject: Addendum, BEERBUG

[From Bill Powers (920307.0900)]

I glossed over the bug leg details a bit in the previous post.
Consider two opposing muscles operating a limb around one joint.
A neural signal shortens a muscle by making its contractile component shorter. A force is developed if the spring compnent is stretched relative to the muscle's shortened state. We can lay out two muscles in a straight line, anchored at the ends (with Xs), to show how a balanced pair of muscles responds to a balanced pair of signals. The center point (marked O) moves toward the side with the larger signal.

```
           large signal   small signal
                 |              |
           X//////////O/  /  /  /  /  /  /X
                      |
                      |
                 position of
                 center, no
                    load
```

If you imagine a pulley at the position of O, and the two springs draped over it, you can convert to angular motion around a pivot. The distance x is then measured in radians.

If an external force is now applied at O, positive to the right, the center will be deflected to the right by an amount depending on the passive spring constant of the muscle, k:

```
      small signal       large signal
                 |       |
     X/  /  /  /  /  /O///////X
                      |
                      |  Fr <--  (restoring force)
                      |  Fa ---> (external force)
                    x0|       x
                      | ----->|
                      |   position of
                      |   center with
                      | deflecting force
                      |
                 position of
                 center, no
                    load
```

If S is the difference between left and right signals, then x0 = cS, where c is the contraction constant (in radians per nervous system unit or NSU). S represents (right - left) signal here and positive distances and forces are measured to the right. The zero-point of x0 is at the midpoint between the anchors.

If k is the spring constant of the combined muscles in force per radian of stretch, then restoring force is

Fr = -k(x - x0), or
Fr = -k(x - cS).

Applied force plus restoring force equals zero, so

Fa - k(x - cS) = 0.

The deflection is then a function of the Signal and the external applied Force:

x = (Fa/k + cS)

This is the "massless" version, which will be reasonably realistic.

If the leg has mass (moment of inertia) and friction, we have a differential equation relating applied force and deflection x:

Fr = k(x - cS) + f(dx/dt) + m(d2x/dt^2) or with an applied force Fa,

Fa - k(x - cS) - f(dx/dt) - m(d2x/dt^2) = 0

where f = coefficient of friction, and "m" is really the moment of inertia about the joint. It actually gets a lot more complex than that, which is why I don't recommend using a completely real physical model without a pretty high-powered mechanicist helping out.

Best,   Bill P

9203B CSGnet

Date: Sun Mar 08, 1992 3:18 am PST
Subject: BEERBUG models

[From Bill Powers (920308.0400)]

Greg Williams, thanks very much for the articles by Cruse (two-part article on "Quantitative model of walking... ", 1980) and Graham ("Simulation of a model for coordination of leg movement...," 1977), from Biological Cybernetics. As I suspected, I have been reinventing a number of wheels, but it's nice to see that I haven't misconstrued the problems that others have been trying to solve. These papers offer models containing many features that I've incorporated into mine -- Cruse even proposes negative feedback control systems for leg position. The main gait-control solutions (as in Beer's model) entail circuits for resetting the legs to a forward position after a complete stride, with the reset from one leg circuit inhibiting and then initiating the reset for the next leg forward (or backward), and with contralateral circuits interacting to make lateral pairs of legs alternate steps.

While there are, as I've been discovering, many circuits that will produce these effects, the least complex methods seem to boil down to some basic logic that's pretty much forced on the designer. The trigger for a forward swing can come from limit-detecting sensors, from comparing leg position sensory signals with upper and lower limit values, and even from comparing the muscle driving signal with upper and lower limit values (this last is the equivalent of a central oscillator). All three methods might be present. Of course the legs must lift during the reset. The speeds of forward and backward swings can be produced by an integrator with variable input currents, the design with which I began. Neither author used my method based on negative feedback control of angular rate of change (Beer used an assumed force-velocity relationship in the legs or body to produce variable speed). I'm not sure that negative velocity feedback is the best solution -- it may be overkill for bugs.

Experimental data seem to require negative feedback control of leg position: for example, the force exerted by a leg rises when the insect drags a weight, ruling out open-loop motor outputs. This was the only control-system experiment done, if I remember right. No data on control of velocity was presented in either paper; velocity might be under control or just produced open loop.

I think it's possible to make a case for a configuration level (position control), a transition level (velocity control), an event level (lift-swing-drop during reset), a relationship level (signal limit detection), a sequence level (propagation of resets from one leg to another) and a logic level (logical conditions on mutual inhibitions among events). No category level that I can see. These are not complete control systems, necessarily -- some may be considered open-loop (I could also be imagining them). The reset event, for example, is stereotyped and there's no need for the bug to sense it or recognize it -- only to produce it. This is very interesting from the standpoint of development or ontogeny. Could control of different levels of variables grow out of an initial capacity to vary them open-loop? In simple enough systems, which limit themselves to simple niches, higher-level variables like reset events really aren't subject to disturbance under normal conditions, and there's no requirement for precision. Something to think about seriously.

Cruse is using the method of modeling in the right way: trying to build a model that will account for observed aspects of behavior. The modeling does have a forcing effect on the data, however. I am still dubious about the notion of a central oscillator, for example, because the data really show a LOT of variation in phase relationships between legs, which a central pattern generator shouldn't produce. I still favor looking for an asynchronous model in which velocity can be smoothly varied between maximum forward and maximum reverse with resets occurring automatically. Maybe both methods are used.

I had occasion to do a behavioral experiment today. We have little brown and gray bugs that come in with the firewood -- they look like "shield bugs" but not quite; about 1 cm long, with long thin curved feelers in front. I saw one today climbing a glass door and applied a few gentle disturbances. This bug, fortunately, was either cold, lethargic, or stupid because it let me fool with it for a while without panicking. I hope they don't carry some deadly disease.

It moved slowly when I barely touched it from behind -- sometimes several touches were required to keep it moving long enough to see its gait. This was not a terrified bug, or else it was terrified into near-paralysis. I couldn't tell whether the reset was propagating forward or backward, but it may have been forward. Resets were not very fast. The legs did NOT reset one after the other in a fixed time, 1-2-3, so the model with a fixed temporal sequence of resets wouldn't work for this bug (unless I was seeing its fastest gait!). Contralateral legs did produce alternate steps: you could see each pair of legs waddling along.

The most interesting phenomena showed up when I applied a disturbance from in front, with the bug stationary or moving. Touching an antenna on one side only, from directly ahead, resulted in immediate swiveling of the front of the body in the other direction, around a pivot near the hind legs, without backing up or with only a small withdrawal. Touching both antennae at once caused a stop, and pushing resulted in a very nice reversal, straight back, similar to the forward stepping pattern but not as coordinated-looking. I could get it to back up only a few steps before a different behavior appeared.

After a couple of head-on pushes, the bug went into a sideways crab-like movement, its body remaining oriented straight into my finger but moving exactly sideways. All three pairs of legs went through steps with resets just as if it were moving forward, but only lateral leg movements were occurring. Contralateral legs still alternated steps, as before. I couldn't see the sequence fore-and-aft -- I was too astonished by the crab gait. When I persisted, the bug turned its body by crabbing the hind legs one way and the forelegs the other way, with the middle legs hardly doing anything. It pivoted about its own center. There is clearly independent control of the leg-pairs.

All this was taking place on a vertical glass surface with the bug moving mostly diagonally upward. When I finished, I realized that this bug (which I was originally going to squash) had become a pet, so I had to scoop it up on a piece of paper and deposit it in the great Outdoors, where it was probably eaten within 5 minutes. But thanks for the game, bug.

I now realize that there are sideways gaits with resets, and that there are several ways to use them to produce turns with radii down to zero. The same

kind of circuit should be usable for crab motion as for forward motion, but using the lateral muscles. I don't know if this is observable in other species, or even other examples of my bugus minimus. One thing I am sure of: in the bug I observed, collision with an obstacle does NOT result in a stereotyped backing and turning, but in clear control of pressure on the antennae by means of perfectly appropriate reversing, turning or crabbing movements, just sufficient to reduce the pressure to zero. I don't know what a more excitable bug would have done.

All of this, plus reading the literature you sent me, reaffirms my view that before a definitive model can be constructed, behavioral experiments have to be done, and specifically experiments aimed at revealing controlled variables. Heroic measures like amputations and deafferentation may tell us something, but I feel that working with intact animals is likely to mislead the least. The details of neural circuitry in these models are not very constrained by knowledge about actual nervous systems -- it's mostly a matter of coming up with a design that will reproduce various features of the behavior. Some of these features -- like how long a leg remains up during the reset phase -- aren't very important, yet it's easy to get hung up on trying to get perfect reproduction. There are too many ways to reproduce such details; I think it's best to try to get the major aspects of the behavior right, and let modifications be added by those who are truly interested in bug neurology for its own sake. There are a jillion reset circuits that would all work the same way. I'd just as soon pick one that works reasonably well and get on to more interesting (higher level) behaviors. When somebody actually traces the real reset circuits, we can just rub out that part of the model and fill in the right circuit. It won't DO anything remarkably different.

One thing this modeling can do: it can help circuit-tracers recognize what they're looking at. So it's good to have several possibilities for the various functions.

Also, I keep remembering the large variability in neuroanatomy from one bug to another of the same species. Looking for THE gait-control circuit may be a wild goose chase. Maybe bugs just reorganize until they have A gait-control system that serves their higher-level purposes. In that case, we might find that ALL the models proposed, that work, will be found in one bug or another, even in one species. I doubt that all these behavioral experiments mentioned by Cruse were done with many individual bugs, to check that they all work the same way in detail. Cruse mentions that some of the observations are contradictory. Maybe that just results from the fact that when you've seen one bug, you have NOT seen them all. If Beer's "genetic algorithm" modeling succeeds in creating functional bugs, it will stop with the first design that meets the criteria -- and it might never produce the same design twice, even though all the designs "do" the same thing.

Best,   Bill P.

Date: Sun Mar 08, 1992 1:52 pm PST
Subject: Re: Powers Reply

[Martin Taylor 920308 16:45]
(Rick Marken 920306 -- I've been incommunicado since Friday)

Rick says:
PCT begins with the observation of a phenomenon -- control.

I take control to be a theoretical construct, not an observation. If it had been an observation, would it have taken Bill's insight to see it? The word "insight" seems appropriate -- control is hidden "in" what we observe, just as are the percepts that are the objects of control.

From my perspective, control would be only another interesting phenomenon, were it not for the fact that control is essential if living organisms are to survive on an evolutionary time scale. It is the necessity of control rather than the "observation" of it that makes it interesting. And once its necessity and its occurrence have been noted, then one can see that it accounts for a lot of other phenomena that might otherwise seem mysterious or unrelated.

Martin

Date: Sun Mar 08, 1992 6:19 pm PST
Subject: Phenomenon of control

[From Bill Powers (920308.1800)]

Martin Taylor (920308.1643) --

>I take control to be a theoretical construct, not an observation. If it
>had been an observation, would it have taken Bill's insight to see it?

It's an observation, but an insight was needed to bring it to attention.

This is the situation observed:

```
       ------------                 -------------        --------
      |            |               |             |      |        |
      |  BEHAVING  |----->ACTION-->|ENVIRONMENTAL|----->|OBSERVED|
      |   SYSTEM   |               | TRANSFORMA- |      |OUTCOME |
      |            |               |    TIONS    |      |        |
       ------------                 --------------       --------
                                                            ^
                                                            |
                                       Link to outcome ---->|
                                                            |
                       VARYING INDEPENDENT DISTURBANCE -->-
```

The observed outcome is maintained stable against the disturbance by variations in the action of the system. The independent disturbance itself is NOT sensed by the behaving system. The effect of the transformed action on the observed outcome is nearly equal and opposite to the effect of the varying independent disturbance on the same outcome.

When the above situation is observed, the behaving system is said to be controlling the observed outcome and the observed outcome is then called a controlled variable. The above diagram is an operational definition of control.

The reason this phenomenon has not been recognized is that "Scientific Method" (as taught in psychology) systematically rejects controlled variables. The action of the behaving system (interpreted as a response) will show a negative correlation with the varying independent disturbance (misinterpreted as a stimulus), to the extent permitted by the environmental transformations. A mistaken inference is involved:

```
             --------<---(correct inference)-------<--------
            |                                               |
       ------------                 -------------        --------
      |            |               |             |      |        |
      |  BEHAVING  |----->ACTION-->|ENVIRONMENTAL|----->|OBSERVED|
      |   SYSTEM   |  ("response") | TRANSFORMA- |   |  |OUTCOME |
      |            |      \         |    TIONS    |  |   |        |
       ---------__-         \       --------------    \  --------
               |\             \<-- (poor correlation    \   ^
        (mistaken \ inference)  \    but possibly         --|<--(high neg
                    \             \   significant)          | correlation,
                      \             \                       | highly sig-
                        \             \                     | nificant)
                       VARYING INDEPENDENT DISTURBANCE -->--
                                 ("stimulus")
```

The negative correlation mentioned above will be maximized only when both action and disturbance are translated into units of effect on the observed outcome. If other measures of the action and the disturbance are used, the obtained correlation between measures of action and disturbance may be either negative or positive (depending on the measurement scale); it is also is likely to be considerably less than perfect.

The controlled outcome will show a low correlation with both the actions of the system and the independent disturbance. The better the control, the lower the correlation. Therefore the observed outcome or controlled variable will be discarded by a statistical analysis because it shows no significant relationship to either "stimuli" or "responses." It will not be recognized as an outcome of behavior. The mistaken inference will be that the varying disturbance is being perceived, because it correlates with a response (and thus satisfies the behaviorist's operational definition of a stimulus). The slanted line in the second diagram shows the mistaken inference. The correct inference (shown by the dashed line at the top of the second diagram) is that the behaving system is sensing the controlled outcome. It can be verified that when this sensing is prevented, control disappears.

Thus control is not a construct, but an observation. To detect control, it isn't necessary to know anything about what is inside the box labelled "Behaving system." It isn't necessary to know what a control system is or how it works in order to know that control is going on. The Test for the Controlled Variable can reveal control by identifying a variable that fits in the position of the "Observed Outcome" above: The Test is simply a way of verifying that the conditions shown in the diagram are present.

Best,   Bill P.

Date: Sun Mar 08, 1992 6:28 pm PST
Subject: Re: Powers Reply

[From Rick Marken (920308)]

First, I'm sorry to hear that Beer is not on CSGNet. But thank you, Gary, for forwarding my post to him. If Beer posts again to CSGNet then I will reply to his address and to CSGNet. But this comment is just going to CSGNet in reply to Martin Taylor (920308 16:45) who says:

>Rick says:
>> PCT begins with the observation of a phenomenon -- control.

>I take control to be a theoretical construct, not an observation. If it had
>been an observation, would it have taken Bill's insight to see it? The
>word "insight" seems appropriate -- control is hidden "in" what we observe,
>just as are the percepts that are the objects of control.

About five years ago (after working with control theory for over five years) I realized why I was having so much trouble getting conventional psychologists interested in control theory. It was because they did not recognize the phenomenon of purposeful behavior -- ie. control. I think Bill tried to make this point in BCP. But it was never really clear to me until I realized that psychology (and all social science) really assumes that behavior is generated output -- NOT CONTROL. Then I realized that control is a phenomenon that can be observed and distinguished from similar apprearing phenonmena. That is the point of my mindreading demo and some of my other demos (like the "findmind" demo). It shows that you can distinguish controlled variables from other, uncontrolled variables WITHOUT ANY THEORY OF CONTROL. The theory does help you understand what to look at -- it's true. But even without the theory you could tell that some variables are under control; disturbances do not have anything close to their expected influence on these variables, and this is usually traceable to systematic opposition to the disturbance by other influences on the variable.

I believe that the recognition that what we call "behavior" IS CONTROL -- ie, that behavioral variables are controlled variables -- is Bill Powers' revolutionary insight. It is not so much a theoretical insight as a phenomenal (in all senses) one. Bill realized that organisms produce consistent results by variable means. So right there a straight through causal model of how the end results are produced is not going to work. He also realized that the variations in the means are exactly what are needed to compensate for other influences on the result besides those of the organism; ie, control is going on. The fact that Bill is trained in physics probably helped him see this -- a stupid psycholgist like me would have been happy to say about a behavior like "take a drink" - "no problem, you just lift the glass and drink". But Bill knew that the forces involved were different each time; so the muscles have to be doing something fairly different each time or there will be no consistency.

Once you know about the possibility of controlled variables then you can start looking for them -- no theory needed. They are like fossils for evolution. People knew how to find fossils and how to classify them before they had any idea that they were telling a story about the history of life. You could probably find out a whole lot of interesting stuff (about people and bugs) without necessarily knowing control theory. All you have to know is that controlling is going on and how to test for controlled variables. The theory will help you determine if you understand HOW these variables are being controlled.

Of course, another of Bill's brilliant insights was the realization that the variables being controlled must be perceptual variables; this is important because it suggests that some of the variables an organism is controlling may not be the same as the variables we identify in our models of physical reality. Controlled variables might be things that a physicist says don't exist -- like, my favorite, honesty -- but can be perceived and must be derived from other variables that are, ultimately, "out there".

So I will stick to my claim that control is a phenonenon. And I will also say that it is important to recognize it as such. That way, we will start getting studies (hopefully) of what we (PCTers) need to know about the most -- what variables organisms control -- and which they do NOT control. As Bill noted in his discussion of the bugs, there is precious little DIRECT information about controlled variables that is available. PCT requires a whole new approach to research -- one geared to investigating the phenomena of control. That is why PCTers have such a hard time dealing with the current literature in psychology (and sociology, biology, zoology, etc). Very little (if any) of the behavioral data is based on an attempt to determine the variables that organisms control. What we (PCTers) need are clever and industrious researchers; once you've got the data the modeling is not really THAT hard. And besides, when you have the data at least you know what your model is supposed to do.

I hope this is reasonably clear. I look forward to hear what you think about it.

Hasta Luego   Rick

Date: Mon Mar 09, 1992 6:14 am PST
Subject: The Curse of Control Theory

[From Gary Cziko 920309.0745]

Bill Powers (920308.1800) noted concerning the phenomenon of control:

>The controlled outcome will show a low correlation with both the actions of
>the system and the independent disturbance. The better the control, the
>lower the correlation. Therefore the observed outcome or controlled
>variable will be discarded by a statistical analysis because it shows no
>significant relationship to either "stimuli" or "responses." It will not be
> recognized as an outcome of behavior.

This, I believe, is the hardest part of understanding control and getting others to see it happening in living organisms. A high correlation (either positive or negative) is one which we have been taught (even since psychology became "scientific") is important. Low ones (close to zero) mean that two variables are unrelated--and yet it's amazing how low the correlations can be and still be considered "significant" and important if they are between a type of stimulus (independent variable) and response (dependent variable).

I wonder if control theory will ever gain widespread understanding because of this absolutely foreign perspective on correlation. I'm not even sure that *I* really understand how the "correct inference" involves taking the path of the close-to-zero correlation. If I didn't keep going back to DEMO2 I think I would be lost to CT as well. And the fact that in no other science (that I know of) uses low correlations to find out what is happening doesn't help matters. Bill, is it the case that even the engineers who create artificial control systems don't need The Test since they know what they want to control and it becomes quickly obvious if they have succeeded or not without calculating the low correlations?

I therefore strongly encourage Bill and any others who have insights on how this can be made understandable to post them to the net. Perhaps the CT old-timers have forgotten how difficult this radical idea is that the important relationship is shown by a near-zero correlation. I wonder how many people on this network have even a basic understanding of this. I think that this is really the CURSE OF CONTROL THEORY (which is, of course, also it's most important insight).--Gary

Date: Mon Mar 09, 1992 6:44 am PST
Subject: language

[From: Bruce Nevin ()]

Bill Powers (920305) --

There are so many hypotheticals in this post, it was a bit difficult to digest. Also, it was unclear whether you were offering a game of "let's you and him fight" or conciliating "Boys! Boys! Don't fight!" Setting all that aside I will focus on the conceptual variables that seem to matter most to you. Please correct my aim if I miss.

> anything common to all languages will be found at the level of
>nonverbal experience, not in language conventions (except as these
>conventions are inherited from other languages).

>If we can trace certain structural constraints to the world of
>nonverbal perception, then they will no longer have to be explained in
>terms of rules relating words as words.

I think any linguist would agree with this, though there is much controvesy about where the boundaries lie, how to determine such boundaries, how to characterize the different kinds of orderliness on either side of such boundaries, and so on--mutually interdependent questions.

There are contributors other than HPCT to linguistic universals, but it seems to me that they are necessarily all mediated by perceptual

control. An example is the set of acoustic properties of the vocal tract that results in favoring certain places of articulation for consonants, which I have sketched a couple of times. In part, the infant encounters utterances in which consonantal bursts, transients, etc. are mimicable only by configuring the tongue so as to constrict the vocal tract in these favored regions; in part, the exploratory self-unfoldment of the control hierarchy finds experientially that in these regions articulatory error makes less acoustic difference than does a similar difference of configuration and effort at other regions of the vocal tract. It is doubtful whether the infant would make this sort of discovery without the prior existence of language (conforming to these constraints) as a conventional social artifact in the environment. Even with an environment providing many experiences of language use, evidence is that infants do not develop requisite control without the motivation of being able to use language to engage others in interpersonal communication and to accomplish personal goals through cooperative social means. Children brought up to age 4 with little human interaction but with their cribs next to a TV that was constantly on were severely impaired in their linguistic and social skills, though presumably exposed to a great deal of very sophisticated use of such skills that happened not to engage them in interpersonal ways.
(Of course, Bruner's LASS has a much more active role than I imply here.)

If you accomplish the aim of accounting for what all languages have in common, and you show that it all comes down to characteristics of the world of nonverbal perception plus fundamentals of physics and chemistry in the environment, like the acoustics of the vocal tract--having reached the state where linguistic universals are trivially deduced from first principles, what would remain? In your hopeful estimation, the conventional aspects of language would be simple and uncontroversial, and the different systems for describing it would converge. I believe that it would remain quite complex. And I believe that using some existing systems for describing language (both aspects together) makes an approach to the derivation of linguistic universals from first principles unlikely.

In particular, I believe that operator grammar shows a simple structure for language--a structure of word dependencies--that is universal and that accords well with perceptual control, plus a more or less complex and arbitrary, institutionalized system of conventions whereby words and word dependencies may be given different shapes in utterances. The principles and some of the patterns of the reduction system are universal, but the detailed reductions and the particular word shapes are not.

You object to what you call the "expansions" of operator grammar as being unnatural and not corresponding to your introspective "feel" for what you are doing when you use language. Part of the problem is some confusion about the status of these changes of form. In the example of analyzing a sentence I quoted from your prior post, I hinted at part of the resolution when I said that a particular expansion need not be carried out, it only had to be available, and that the expanded and reduced forms were alternative forms for the same words and word dependencies. Another part of the problem concerns the difficulty of introspection and what is available to awareness, compounded by the fact that you are using the thing you are analyzing and analyzing it even as you use it.

I'll try to address both aspects.

Martin Taylor (920304 17:00) --
Rick Marken (920305) --

>consciousness of means vs ends depending on level
>of disturbance

This is relevant to the discussion of how "natural" a model of language control appears to us as we use language.

>if the actions that allow high-
>level control are easy, then what subjects see themselves as doing is what
>we would call satisfying the high-level reference. But if the lower level
>control structure is disturbed or not well structured (the actions are more
>difficult), then people see themselves as "doing" the low-level things.

The things that are difficult in language control, and which therefore obtrude themself on conscious awareness, are mostly larger discourse structures across series of sentences. Even when a single sentence must be recast, it is typically due to relations in a larger discourse context, and involves reduction to one complex sentence constructions that could also be articulated as two or more sentences. This is the problem of parcelling global, nonlinear word/percept dependencies out into a linear sequence of linearized dependencies constituting sentences. In my master's thesis in 1969 I called this periphrasis as distinct from a paraphrase process within the scope of a sentence. These paraphrase processes seldom rise to awareness.

Lower-level changes of word shape within these paraphrase processes--what linguists call morphophonemic alternations--arise to consciousness for the average language user only when they become socially marked as shibboleths of region, community, or social class. Things like "ain't" and "She don't know no better." These constitute the tiniest, though most visible, fraction of what is going on.

The reductions of operator grammar account for sentence-paraphrase processes. They account at present only for those aspects of discourse periphrasis that are closest to sentence paraphrase, by reductions of conjoined sentences and reductions to pronouns and other referentials. It is important to understand that the reductions include and are no different from morphophonemic alternations of word shape. Let's look at that.

We feel that geese is the same meaning/word as goose plus the same plural meaning/element as the -s of picnics, the -es (that is, -iz) of foxes, the -en of children, and the zero of fish (alongside fishes) or of series. We say that went is the same meaning/word as go plus the same -t that is found in swept, which is none other than the past-tense meaning/suffix that also takes the shape of -ed in braided, the -t of swept, the vowel difference of break/broke, the zero of . . . well you get the picture.

Some changes of form are optional.

John came with Alice, but Alice didn't leave with John, Alice left with Frank.

John came with Alice, but Alice didn't leave with John, she left with Frank.

John came with Alice, but she didn't leave with John, Alice left with Frank.

John came with Alice, but she didn't leave with John, she left with Frank.

John came with Alice, but Alice didn't leave with him, Alice left with Frank.

(etc.)

The differences here are differences of emphasis, not of meaning.

The words in a sentence may be given different forms when combined in specifiable ways with particular other words. That is what the reductions of operator grammar are about.

You, Bill, want to say that it is the *meanings* that are given different word-forms under different conditions. But the conditions ("environments" as linguists say) are not specifiable in terms of other meanings, but only in terms of other words representing meanings.

Its worse than that. You may recall that I asked some time back how an elementary control system (ECS) could control for two input signals being repetitious or redundant with respect to each other ("the same"). You can say she instead of Alice in the above examples only if both words refer to the same individual. Hearing she amounts to hearing an assertion that the individual who arrived is the same as the individual who left. The reduction to she is one form in which that metalinguistic assertion can be uttered. Since I don't know how an ECS can control for sameness (same reference) of its own inputs or outputs, I think some other ECS controls for an *assertion* of sameness. In other words, absent an answer to my question (above) it appears to me that perception of sameness requires metalanguage (as part of language) or a precursor very much like it in prelinguistic control of "metasymbol" perceptions about symbol perceptions (control of one sort of perception as a symbol for another).

What could that be? Possibly the first step toward language is the ability to perceive repetition (one instance or token of a category and another instance or token of that category). Some imagined perception is taken as a symbol for any instance of the category. Pictographs, hieroglyphics, and ideograms work like this. Rebuses do also, but in an ad hoc way that has not be institutionalized. Withal, we must be careful to avoid identifying the evolution of writing systems too closely with the evolution of language. However suggestive the parallels may be, they are still separated by a great span of evolutionary time.

A next step toward language must be the perception that two symbols (category perceptions) refer to the same individual instance or token. This "sameness" perception relative to the category perceptions is not a category perception, but rather is about the category perceptions in precisely the same way that a linguistic utterance is about the perceptions to which it refers. (And indeed, by that relation it may appear to create the act of reference, and the relation of reference between a category perception and the perception of individual token of the category, but that is rugged epistemological terrain that I can only look at for now, not enter upon.) With these two evolutionary steps, category and metacategory, you have the first requisites for language. (Metacategory has nothing to do with the question of categories of categories we have discussed in the past.)

This metalanguage referring to the words of language is itself a part of the language, using a subset of its words and a limited portion of its syntax. Metalanguage assertions are almost always reduced to morphemes like pronouns and articles. They are thereby made especially difficult to notice.

Date: Mon Mar 09, 1992 9:13 am PST
Subject: Re: The Curse of Control Theory

[Martin Taylor 920309 11:00]
(Gary Cziko 920309.0745)

Gary says that the key to seeing control is the finding of zero or near zero correlation. Inasmuch as the correlation between almost any two variables in the universe is very near zero, that test would lead to the conclusion that almost everything is linked by control.

(Bill Powers and Rick Marken 920308)

Bill and Rick both assert that control is a directly observable phenomenon. I had said I thought it was a theoretical construct rather than a direct observation, by which I meant something like "blueness" or "tastiness." I think we are likely to get into an unnecessary war of words, here. To Bill I would suggest that we have long agreed that the discovery of a possibly controlled percept is often difficult and requires insight before the Test can be applied. Even then, it is usually not easy to determine what disturbances are occurring, and it is only for very low levels that the outside observer (experimenter) can determine with precision that the putative controlled variable in the subject's perception is close to the one "observed" by the experimenter.

To Rick, I would point out that William James lived some time before Bill Powers. One may credit Bill with a lot, but not with the discovery that people achieve one purpose by variable means. And I doubt that James would have taken it as an original observation, when he said: "Provided the same conclusion be reached, the means may be as mutable as we like, for the "meaning" of the stream of thought will be the same. What difference does it make what the means are? "Qu'importe le flacon, pourvu qu'on ait l'ivresse.""

If, as an experimenter, one can presume some pattern in the mutually observable environment represents a perceptual variable being controlled by the subject, then one can attempt to disturb that pattern and see whether the subject acts so that the pattern is restored or maintained. The pattern will show little correlation with the experimenter's disturbances or with what the experimenter observes of the subject's actions. If the experimenter happened to be correct that what she did would have disturbed the pattern if the subject had not been there, then there is evidence that the subject is controlling. The presumption that the experimenter would have disturbed the pattern is just that, a presumption. It is not an observation, because it didn't happen. Explaining why things do not happen is trickier than providing rationales for why they do happen. The failure of a presumed "cause" is easier to justify as that it was not a cause than as that an exactly countervailing cause was applied at the same time. I think this is at the root of the communication difficulty with cause-effect psychologists. Causes have effects, and PCT is supported when what should be causes are observed to have little or no effects.

In PCT terms, one can model a subject and an experimenter seeking the subject's perceptually controlled variables as being two control systems with conflicting references. The experimenter's reference is that the presumed percept of the subject should be altered, and the subject's reference is that it should not be. This is an interesting contrast to the description of communication, in which the two parties are assumed to have common goals (that each should be satisfied that the communication has happened as the other desired). And that raises interesting questions about ethics in psychological experimentation.

Martin

Date: Mon Mar 09, 1992 10:56 am PST
Subject: Re: The Curse of Control Theory

[from Gary Cziko 920309.1220]

Martin Taylor (920309 11:00) responds:

>Gary says that the key to seeing control is the finding of zero or near zero
>correlation. Inasmuch as the correlation between almost any two variables
>in the universe is very near zero, that test would lead to the conclusion
>that almost everything is linked by control.

You see, I was right--the curse strikes again.

What Martin's comment points out seems to be that the low correlation between the controlled outcome and the actions of the system is a special type of low correlation. It is not a low correlation given by an amorphous blob of points on a scattergram indicating that any value of x can occur with any variable of y. Instead, it is a low correlation given by lots of different values of x (action of the system and/or independent disturbances) occuring with only a small range of values of y (controlled outcome).

So a low correlation is not enough. It is a special type of low correlation. And it would disappear of the perceptual link between controlled outcome and behaving system were eliminated.

I hope Bill and/or Rick will have something to add to this.--Gary

Date: Mon Mar 09, 1992 11:12 am PST
Subject: anatomy of a corpse

[From: Bruce Nevin (Mon 92019 12:37:25)]

I'll pick up on this one as a test of comprehension. Please disclose holes.

Low correlation is not a diagnostic for control. As Martin says, there is low correlation between too many things in the universe. Things that exhibit low correlation with one another cannot be relevant for scientific understanding. That includes particularly the relation between observed outcome, "stimuli" and "responses" (actions).

It seems to me that Bill's point is rather like the old Sufi teaching story of the man looking for his keys under the street lamp. I think probably we all know it. A friend
stops to help. After a while he gives up.

> Are you sure you lost them here?
> No, I lost them over there.
> Then why on earth are you looking here?
> There's more light here. It's all dark over there!

The correlation that matters takes us into a place where there is less light--the correlation of an outcome with a purpose. The outcome is deep in murky shadow because it comprises selected aspects of a situation, leaving out other aspects that are irrelevant. What is the criterion of relevance? The purpose is completely in darkness because it is a memory of perceptions that are present in the situation. This memory is internally maintained in the "subject" of the experiment, not accessible to the experimentor. Nothing "there" in the experimental situation to observe!

But when you observe a series of different disturbances being corrected you get some light on the purpose or goal. In the changing situation something emerges as invariant. (It's more or less nearly invariant, depending on the gain, and the invariant might itself involve a perception of change, sequence, etc., so the term "invariance" can be a stumbling block.)

This invariant, the observable outcome, is not directly observable as a first-order observation. It is a second-order observation, observable only in context of certain expectations derived from hierarchical perceptual control theory.

But even this observation of invariance in a shifting situation of the "subject" and its environment is not the thing that correlates with remembered perceptions within the "subject." It *reflects* the internally-maintained goal. Only some aspects are relevant, and those relevant aspects might not even be noticed from the point of view of the observer. The investigator must shift perspective to recognize what perceptions of the observed outcome are available to the "subject."

When you have some idea what the controlled outcome is you can apply the Test. But the Test doesn't make much sense until you have identified an invariant in the situation from the point of view of the controlling organism.

So the unfamiliar steps include isolating an "invariant" outcome in the organism/environement system, selecting those aspects of the outcome that are relevant from the point of view of the organism, and applying the Test to verify that your guess as to what is being controlled is correct.

This is all complicated if the organism starts controlling for something else for whatever reason.

Low correlation is relevant only for folks looking around the lamppost. It's a way of telling them that the keys aren't there.

Bruce

Date: Mon Mar 09, 1992 11:12 am PST
Subject: obliterating variation

[From: Bruce Nevin (Mon 92019 13:20:24)]

Another neglected wildcard involves genetic variation. If we are talking about genetically innate mechanisms, there is little reason to suppose that human beings are genetically homogeneous. Alletic variation in humans is at least 7%, according to Lieberman. There is considerable difference in speech perception in human beings that seems to be correlated with genetic differences.

Lieberman _The biology and evolution of language_ (206) draws an analogy to differences in respiratory function. There are four different types found in the population, arguing for four underlying genetically transmitted mechanisms for oxygen transfer, "despite a hundred million years or so of evolution involving the respiratory system. By studying the pattern of variation, the physiologist can, however, predict the [respiratory response] behavior of an individual once he knows what group the individual falls into.

> Consider instead the treatment of these physiologic data that would occur if they were cast into the competence-performance model that derives from the linguistic theories of Saussure and Chomsky. The linguist-physiologist would attempt to find a single biological mechanism that reflects the underlying "respiratory competence" of all 33 subjects. The differences in behavior would, to the linguist-physiologist, reflect the presence of mysterious "performance" factors outside the domain of linguistic physiology. Faced with these data, the linguist-physiologist might derive an "average" function that accounts for the behavior of none of the subjects. Most likely the linguist-physiologist would have to conclude that since no single competence function could be derived that accounted for the behavior of all the subjects, the problem was outside the proper domain of linguistic physiology.

It is unclear how much of this argument transfers from phonology to syntax and semantics. I suspect that there is a great deal of variance in the population that gets swamped in the rush for universals. Perceptual control in all cases.

-=+=-

On another front, selective adaptation experiments trade on the notion that parallel "feature detectors" compete for interpretation of sensory input.

> Listeners are first asked to identify speech stimuli that differ with respect to an acoustic continuum, such as the second formant frequency transitions that will cue the consonantal place of articulation. This initial session establishes the phonetic boundaries along this continuum. In a later listening session the same subjects first listen to several trials of an *adapting* stimulus, such as a stimulus that has been identified as a [ba], followed by five or so stimuli drawn from the test continuum. The outcome of selective adaptation is that a phonetic boundary shifts. If listeners adapt to a [ba], stimuli that were previously identified as [ba]'s that were near the [ba]-[da] boundary . . . will now be identified as [da]'s. The theoretical interpretation of selective adaptation is as follows. . . . If speec signals are hypothetically identified by triggering feature detectors, then the location of a phonetic boundary represents a point along a continuum where two detectors are responding with equal strength. Fatiguing one of these detectors by presenting a number of trials would reduce the response of the feature detector that responds to the adapting stimulus. This would yield a change in the balance point along the continuum where both feature detectors respond with equal strength. A shift in the phonetic boundary toward the adapting stimulus thus is consistent with this theory and has been interpreted as evidence for feature detectors for speech perception. (ibid. 184)

Fatiguing?! One would then expect other concomitants of fatigue. Continued repetition of the "adapting stimulus" should fatigue the relevant feature detector right off the map so that only [da] is heard eventually, even when a [ba] stimulus is presented. One would expect diminished capacity to hear a certain feature after exposure to an alliterative passage.

It seems to me more plausible that the "adapting stimulus" is used by the control system to establish a reference signal appropriate for a particular speaker. After several variations with little or no variation, the hearer "expects" that anything phonetically different is categorically different (a [d] instead of a [b]). I wonder what happens if the listener is presented with a series of [ba]s and [da]s interspersed, setting consistent reference values for both, and then is presented with the continuum, with the option of saying "neither."

Lunch is over. Got to run to a meeting.

Bruce

Date: Mon Mar 09, 1992 1:58 pm PST
Subject: Skinner on Behavior as Fluid

[from Gary Cziko 920309.1545]

A quote from Skinner which I just came across which I think CSGnetters will find of interest:

"As it stands now, I'm not sure that 'response' is a very useful concept. Behavior is very fluid; it isn't made up of little responses packed together. I hope I will live to see a formulation which takes this fluidity into account."

One could argue that Skinner lived but didn't see.

Source:

Evans, R. I. (1968). _B. F. Skinner: The man and his ideas_. New York: Dutton.

Gary A. Cziko

Date: Mon Mar 09, 1992 3:12 pm PST
Subject: Re: obliterating variation

[Martin Taylor 920309 16:30]
(Bruce Nevin 92019 13:20:24) I think 920309, but that's what the heading says.
>
>
>It seems to me more plausible that the "adapting stimulus" is used by
>the control system to establish a reference signal appropriate for a
>particular speaker. After several variations with little or no
>variation, the hearer "expects" that anything phonetically different is
>categorically different (a [d] instead of a [b]). I wonder what happens
>if the listener is presented with a series of [ba]s and [da]s
>interspersed, setting consistent reference values for both, and then is
>presented with the continuum, with the option of saying "neither."
>

I quite agree with Bruce on this. I have a typescript that I intended to turn into a paper some time in the 1970s on the topic, from which I will quote some passages at the end of this posting. I simply do not believe that any of the perceptual effects attributed to "fatigue" are properly attributed. This applies to figural aftereffects, of which the shift of the phonetic category boundary is probably an instance. It applies specifically to reversing figures, which are often attributed to the fatigue of one percept with the consequent appearance of the other when the fatigue progresses far enough. The timings of the changes of percept are quite inconsistent with a fatigue interpretation and are consistent in detail with a random walk of a small stable number of "detectors." In the case of phoneme detectors it used to be fashionable to say that they could be detected by fatiguing them selectively. I disputed this interpretation of the data, and asserted that any "algorithmic" detector based on subjective probability would reach the same results.

In the quote that follows, the phoneme boundary is between /j/ and /d/, which differ in the artificial stimuli by the length of the noise burst that follows the stop release. After an introductory description of the effect--repeated presentation of /j/ moves the boundary j-ward--the text continues:

"This effect almost always occurs when such an experiment is run, and its occurrence is taken as evidence for the existence of feature detectors for the phoneme of characteristic in question. It is not. Any sensible algoritmic classifier that takes some notice of recent history will do the same.

"In order to see why an algorithmic classifier would shift the transition region toward the 'fatigued' phoneme in a selective adaptation experiment, consider the distribution of occurrences of the characteristic (e.g. burst length) in examples of phonemes from natural speech. Some examples of /j/ will have short noise bursts, some examples of /d/ longish ones. In the experiment, the subject must choose on the basis of the length of the noise burst whether the sound is /j/ or /d/. If he has a feature detector, he will select the one whose feature detector gives the higher output. If he detects on an algorithmic basis, his response will be determined by which phoneme is more likely, or to which the given sound is more similar if it falls outside the range of either real phoneme. The situation is as shown in Figure 1 [Picture of sharply peaked /d/ distribution of burst lengths on left, overlapping wide, low /j/ distribution of burst lengths to its right]. If there exists a feature detector, it must be more broadly tuned for /j/ than for /d/, since the range of natural burst durations is greater for /j/ than for /d/. Either feature detector probably has some output for burst durations beyond the natural range of "its" phoneme. If the detector is algorithmic, it must operate on the basis of some kind of "similarity" index related to the probability that a particular phoneme could be represented by the sound in question. This simiarity index probably looks like, but might be broader than the occurrence probability distribution. If the index is likelihood, as it would be for a Bayesian classifier, then it is identical to the occurrence distribution.

The distributions of Figure 1 apply to normal speech. In the experiment, the distributions are quite different. One specific example of the phoneme is presented over and over again. The variance of the occurrence distribution of this phoneme is drastically reduced, by an amount that depends on how sensitive the system is to long-term trends and how sensitive to local context. Considering both local context and history, the distributions (assuming /j/ is being "fatigued") become something like those for Figure 2a [like Figure 1, but now the /j/ distribution is sharply peaked well to the right of the /d/ distribution]. The /d/ distribution is unaffected, but the /j/ distribution is reduced except for the central peak that corresponds to the distribution of experimental presentations.

The responses of the models change, as well. [discussion of feature detector fatigue response omitted]. The algorithmic classifier's index changes to match or to track the changes in the occurrence distributions. It, in effect, says "Now /j/ never gets as short as it used to do, so I know that middle-length bursts that used to be /j/ are now /d/, since the /d/ distribution still hasn't changed." [...] The effect on the experiment is that both the feature detector and the algorithmic detector give the same result, and the standard selective adaptation procedure cannot distinguish them.

[The text continues by proposing that the experiment be modified to retain the natural distribution of burst length for the "fatigued" phoneme during the "fatigue" process, or to increase the natural range of burst duration. In the former case, there should be no change in the perceptual boundary if an algorithmic classifier is used. In the latter, the result depends on the degree to which the natural ranges of the two phonemes overlap. If the overlap is small, the wider range will not much affect which phoneme is perceived during the "fatiguing" trials, and the boundary should move in the direction opposite to the normal finding.]

I think that the same analysis could easily be applied to Bruce's proposal. If we use the term "reference level" as Bruce does, the analysis means that the continuum of perception levels is being re-scaled, by being expanded in the neighbourhood of each presented phoneme. I find it interesting that I used exactly this principle to explain the consistent errors people make in reproducing figures, in my 1960 thesis (Effect of anchoring and distance perception on the reproduction of forms. Perceptual and Motor Skills, 1961, 12, 203-230), and later to account for more orthodox figural after-effects (Figural after-effects: a psychophysical theory of the displacement effect, Canadian J. Psychol., 1962, 16, 247-277). You might like to look at this latter one, because it develops the idea of subjective probabilities in perception better than I was able to do in our discussions before Xmas.

Incidentally, though I never measured correlations, I think from looking at the figures that the figural after-effect paper might satisfy Tom Bourbon's request for an instance in which subjective probability considerations produce precise predictions of experimental data. Indeed, in one later case, the theory provided a good fit to the data using no (zero) degrees of freedom (Psychol Review, 1963, 70, 357-360) even though the variable affecting the size of the effect was not one that had been used in previous one- or two-degree-of-freedom data fits.

That's pretty much a side issue, I suppose, but it has little hooks into several previous discussions.

Martin

Date: Mon Mar 09, 1992 3:34 pm PST
Subject: Re: The Curse of Control Theory

[From Rick Marken (920903)]

Martin Taylor (920309 11:00) says:

>To Rick, I would point out that William James lived some time before Bill
>Powers. One may credit Bill with a lot, but not with the discovery that
>people achieve one purpose by variable means. And I doubt that James would
>have taken it as an original observation

Yes, James knew purpose. I might add that Tolman did too, And McDougall apparently talked about it knowledgeably. So I will refine my claim. Powers' contribution was to quantify the phenomenon of control and show that it applied at all levels of behavioral organization -- from muscle tensions on up. He also presented a quantitative model of the phenomenon.

I do think that James did understand the nature of the phenomenon of control. But because he lacked a model and a way to quantitatively demonstrate purposeful phenomena he never made much of an impact on the study of purpose. James wrote eloquently about purpose in the first chapter of his "Principles.." but the concept is basically abandoned in the rest of the text, and in psycholgy from that time on, I might add. Quite a different story in Powers' test BCP.

Gary Cziko (920309.1220) says:

>So a low correlation is not enough. It is a special type of low
>correlation. And it would disappear if the perceptual link between
>controlled outcome and behaving system were eliminated.

>I hope Bill and/or Rick will have something to add to this.--Gary

Yes, a low correlation is not enough. The special type you mention is what you would expect if the reference level of the controlled variable is fixed. But if it is not then you might see something other than a horizontal line on an scatter plot -- it could even be a circular cloud. What is needed besides a low correlation is knowledge that there WOULD be a HIGH correlation between disturbance and controlled variable if the variable is NOT controlled. What has been forgotten in this discussion of testing for controlled variables is that the researcher chooses the disturbance because s/he knows what its effect will be if there is NO control. In my mindreading demo I know exactly what the effect of the disturbance would be if a number on the screen is NOT controlled. So a different relationship (such as a low correlation) suggests control - NOT because the correlation is low (it will be fairly low for all the numbers) but because it is not what is EXPECTED. So, in order to test for control you must have a pretty good model of how variables interrelate when there is NO control. And even when you do get results that are unexpected, based on the non-control models, you can make yourself more certain that control is occuring by watching it break down (as Gary noted) when you prevent the system identified as the controller from perceiving the hypothetical controlled variable. This can also be shown with the mindreading demo -- when the subject shuts his/her eyes the effect of the disturbance to the controlled number is once again just what is EXPECTED; control has disappeared.

Does anyone have a spare million bucks for me so I can do this all the time?

Ah well, back to work.

Rick M.

Date: Mon Mar 09, 1992 3:38 pm PST
Subject: CLOSED LOOP PERMISSIONS

From Greg Williams (920309)

I'm beginning to think about the next CLOSED LOOP quarterly compilation of CSGnet conversations (due out around April 15). I'm hoping that some of the netters who have been actively participating BUT have not given me explicit permission to use their posts in CLOSED LOOP (for example, Avery Andrews, Oded Maler, and Martin Taylor, among others -- yes, I'm pointing at you!) will complete and sign the following form, and paper-mail it to me as soon as possible: Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328 U.S.A.

Giving permission via the form won't GUARANTEE that some of your contributions to CSGnet will be immortalized in CLOSED LOOP, but it will raise the probability above zero. What goes in CLOSED LOOP is copy-edited and cut to fit by me, but I try to avoid altering meanings, to the extent I'm able (note the form's escape clause if I screw up; also, errata and clarifications could be included in subsequent issues of CLOSED LOOP if warranted). If anyone wants to haggle over fine points of permission-giving, post directly to me or phone me at 606-332-7606.

Honest, folks, this is for a good cause -- maybe Bill P. and/or Gary C. can put in a good word for the CLOSED LOOP endeavor....

TO GREG WILLIAMS:

YOU HAVE MY PERMISSION TO USE EXCERPTS FROM MY POSTS ON CSGNET IN "CLOSED LOOP." I RETAIN ALL COPYRIGHTS TO MY POSTS, AND YOU WILL INDICATE THAT FACT BY INCLUDING A LEGAL COPYRIGHT NOTICE IN "CLOSED LOOP" FOR EACH EXCERPT FROM MY POSTS. I MAY CANCEL PERMISSION (NON-RETROACTIVELY) WITH REGARD TO ANY PORTION OF MY POSTS BY GIVING YOU SIX WEEKS' NOTICE.

SIGNED ____________________________________

DATE ______________________________________

PLEASE PRINT:

NAME ______________________________________

ADDRESS ___________________________________

____

____

____

THANKS,

Greg

P.S. No, I don't plan to include Beer's bug as a topic in the next CLOSED LOOP!

Date: Tue Mar 10, 1992 6:59 am PST
Subject: Re: obliterating variation

[From Oded Maler 921003]:

>[From: Bruce Nevin (Mon 92019 13:20:24)]

>It is unclear how much of this argument transfers from phonology to
>syntax and semantics.

I believe that the case of syntax corresponds roughly to the variations in the types of motor programs that control gait, etc.

Date: Tue Mar 10, 1992 7:00 am PST
Subject: genetic variation; motor programs and syntax

Chomsky has mentioned the possiblity of genetic variation in UG, but it's certainly not on anyone's priority list. Except that there is a family in Montreal, being studied by someone called Myra Gopnik, who apparantly lack the ability to manage grammatical features such as tense and number, and are currently the focus of considerable interest.

The idea of motor programs as a basis for syntax strikes me as quite intriguing, but in a preliminary unguided foray into some of the literature I did not get very far with it.

Avery.Andrews@anu.edu.au

Date: Tue Mar 10, 1992 7:02 am PST
Subject: curses

[From: Bruce Nevin (039210.0828)]
The header "anatomy of a corpse" should have read "anatomy of a curse."
Maybe the mummy's curse intervened.

Martin 3/9 --

Very interesting corroboration. I'll put those refs on my list for my next expedition to a university library. Thanks!

Oded 3/10 --

Motor programs I had thought of as a basis for metaphor and analogy, along with configuration, sequence, and all sorts of other perceptions. Are you saying that, say, an infant's control systems established for program-level control of movement might be replicated and adapted (as a working interconnected system of ECSs that works) to handle language? Forming a kind of template? Can you unpack this a bit? --It's too gnomic for me.

Is that what you were getting at, Avery?

Bruce

Date: Tue Mar 10, 1992 12:03 pm PST
Subject: Blindfolding - RKC

[From Robert K. Clark (491-2499)]

Bill Powers & Bruce Nevin

Sorry to be slow to respond. Other activities have had priority. Hope to improve as I get some of these responsibilities better organized.

Powers comment about "blindfolding" the people at the dinner table makes the "social control system disappear." But this only applies if the person "cares" what the others think of him. Key questions, I suggest are: "Where" is the control system? What is the nature of the "perceived variable(s)" that is (are) controlled? These answers, and perhaps others, are needed to define the System.

If a "Social Control System" exists, it seems to me that a single controllable variable should be identifiable. Composed, undoubtedly of a combination of many (lower order?) variables. Interrupting some of the channels carrying these variables may be expected to change an individual's activities. But WHERE is this System? and what and where are its Reference Levels?

For a different example of a Social Control System you might consider my discussion of a temperature control system when modern thermostats had not been invented. (See 'STATS VS CRUISE CONTROL, posted Jan 31) That discussion was intended to point out that familiar situations can be regarded as involving a combination of control systems. Such a shift of viewpoint can be quite interesting.

I plan to offer some other ways of analyzing hierarchical arrays of control systems. These will generally resemble our original concepts, but with modifications that I find useful.

Bruce Nevin -- blindfolding breaks the loop

Generally, I agree with you. Several questions are raised by your suggestion, "we don't need all the complexity of hierarchical control to model a human being conforming to a social norm . . . a person seems to emulate an ECS ("elementary"?) controlling a perceptual signal." Perhaps, but how is that "perceptual signal" generated? And how is the set of inputs composing that same "perceptual signal" selected over time? I find the hierarchical structure very helpful in interpreting such complex situations.

Regards, Bob Clark

Date: Tue Mar 10, 1992 12:08 pm PST
Subject: motor programs, caused output, the curse

[From Rick Marken (920310)]

Oded Maler (921003) says:

>I believe that the case of syntax corresponds roughly to the variations
>in the types of motor programs that control gait, etc.

and Avery Andrews says:

>The idea of motor programs as a basis for syntax strikes me as quite
>intriguing, but in a preliminary unguided foray into some of the

>literature I did not get very far with it.

and Martin Taylor (920309 16:30) says:

>If
>he detects on an algorithmic basis, his response will be determined by which
>phoneme is more likely, or to which the given sound is more similar if it
>falls outside the range of either real phoneme.

Motor programs? Responses determined by input probability?

Gary, help, the curse of PCT has taken over CSGNet.

Anybody want to talk controlled variables, maybe????

Perplexed,

Rick

Date: Tue Mar 10, 1992 1:23 pm PST
Subject: Re: blindfolding

[From: Bruce Nevin (Tue 920110 14:18:16)]

Robert K. Clark (Tue, 10 Mar 1992 15:24:00 GMT) --

I couldn't find the posts to which you were responding (if you include a date it helps), and haven't time for more than a quick gasp of air.

Yes, ECS is a recent acronym for "elementary control system," a term offered by Martin Taylor some months ago that we seem to have found useful for distinguishing a minimal neural black box (comparator with I/O functions for perceptual input, reference input, and error or delta output) from a hierarchical control system as a whole.

A key question for social "control" it seems to me is "how can one person model her behavior on that of another?" Something close to that is how social norms and conventions are learned, shared, and used.
The degree to which one follows or violates a norm for X kind of person communicates to others something of what kind of person one is (or purports to be).

In this, it is as though there were an input reference signal for the given social norm or convention. You ask "How is this perceptual signal generated? And how is the set of inputs composing that same perceptual signal selected over time?" I ask: how can one person model his behavior on that of others? If we understand that in a clear and precise way, I believe we thereby have clear and precise answers to your questions.

Bruce
bn@bbn.com

Date: Tue Mar 10, 1992 1:48 pm PST
Subject: Re: motor programs, caused output, the curse

[Martin Taylor 920310]
(Rick Marken 920310)
>
>
>and Martin Taylor (920309 16:30) says:
>
>>If
>>he detects on an algorithmic basis, his response will be determined by which
>>phoneme is more likely, or to which the given sound is more similar if it
>>falls outside the range of either real phoneme.
>
>Motor programs? Responses determined by input probability?
>
>Gary, help, the curse of PCT has taken over CSGNet.
>
>Anybody want to talk controlled variables, maybe????
>
B'ain't nobody here but us percepts.

I was quoting something I wrote about 15 years ago in the passage you repeated. But nevertheless, even though the viewpoint may change, it is still necessary to have grounds for categorization before making the category judgment. Those grounds do not change whether you are dealing with the control of a percept or simply the part of the system between the sensors and the perceptual input to a comparator.

In an ordinary psychophysical experiment, I think the main controlled percept is that of experimenter satisfaction. That being stable, the experimenter is able to induce responses from stimuli in ways we have often discussed. One of them is to provide a stimulus pattern and ask whether it represents class X or class Y, and to express increasing satisfaction as the agreement between the subject's claim of X or Y agrees with the experimenter's opinion of whether X or Y was presented (in other words, the subject does not perceive high experimenter satisfaction by goofing off and answering X and Y independently of what was presented).

I don't think you can simply throw away all of perceptual psychophysics because the subject does not control what the experimenter presents. The controlled percept is at a higher level, and the experimenter knows it. If you have ever watched a naive subject in a psychophysical study, you will know how apologetic they can be when they finish a run. They say something like "I'm so sorry, but even though I tried, I'm sure I got lots of them wrong." It is very hard to convince them that they can satisfy you by doing their best, rather than by succeeding on every trial.

I would now only partially disavow the phrase "response determined by input probability," in that it can be valid if there is some fixed reference against which the percept is known to be compared. (Really there has to be a whole set of references, but many of these are of the same class as Bruce has been dealing with in talking about the socialization of linguistic norms. They are tacitly assumed between experimenter and subject). At the time I wrote the passage I quoted, none of those caveats would have occurred to me. Now they would, but I would consider them part of the experimental environment that justifies the treatment of the data as if the stimuli did determine the responses. When the reference set deviates from what the experimenter hopes it is, the effects are usually pretty obvious. In psychophysical studies, there are usually lots of checks put in for that kind of problem, even if the experimenters would not have described it in those terms.

Does that take the curse off, a teensy bit?

Next--habit patterns and motor programs?

Martin

Date: Tue Mar 10, 1992 4:23 pm PST
Subject: language, motor programs

Re Bruce Nevin (039210.0828)

> Is that what you were getting at, Avery?

Yes. The idea that the pre-adapation for language is guessing what people are up do by watching what they are doing. Given this, one can convey intentions by miming, from which arises manual sign language. Then one has to get from that to spoken language, a step which strikes me as very mysterious.

Needless to say, the term `motor program' should be interpreted very loosely, along the lines of a collection of ECS's such that the one having an error-signal of (approximately) zero is the precondition for the next one doing anything at all (as in my posting about getting a beer).

Avery.Andrews@anu.edu.au

Date: Tue Mar 10, 1992 5:38 pm PST
Subject: Low correlations

[From Bill Powers (920310.1700)]

Look at my address: I now have my own logon at Fort Lewis College, and my new name is powers_w. Legal at last!

Mark Olsen and others interested: back issues of Closed Loop are available for $5 each from

CSG Press   10209 N. 56th St.   Scottsdale, AZ 85253

That's Ed Ford.

To inquiries about my giving seminars at various places: not this year. I have too many committments. I suggest contacting Tom Bourbon, Rick Marken, Kent McClelland, etc.

Martin Taylor (920309), Bruce Nevin (920310) --

You guys are going all around the point here. I have a distinct feeling that you're avoiding it. On the other hand, there may be some critical fact that I don't seem to be communicating, so let's try again:

Here's my diagram of the operational definition of a controlled variable, with those "low correlations" shown:

```
                              ------- low correlation ------
                             |                                |
       ------------          |      -------------        --------
      |            |         |     |             |      |        |
      |  BEHAVING  |----->ACTION-->|ENVIRONMENTAL|----->|OBSERVED|
      |   SYSTEM   |               | TRANSFORMA- |      |OUTCOME |
      |            |               |    TIONS    |      |        |
       ------------                 --------------       --------
                                                         |   ^
                                       --low correlation--   |
                                      |                      |
                                      |                      |
                       VARYING INDEPENDENT DISTURBANCE -->--
```

We have here a situation in which the action of a behaving system affects an observed outcome, and an independent disturbance also affects the same outcome. Ordinarily, we would expect action and disturbance to correlate with the outcome. If action and disturbance were unrelated, we would expect the variance of the outcome to be the sum of the variances of action and disturbance -- that is, in fact, a definition of the statistical test for NO control.

When there is control, the variance of the observed outcome is significantly less than the sum of the variances of the action and the disturbance. When the reference level is reasonably constant, the variance of the observed outcome approaches zero, when the variances of both action and disturbance are large. The only way for this to be possible is for the effect of the action, transformed, to be systematically opposed to the effect of the varying disturbance.

So what is meant by a "low" correlation is one that is less than the expected correlation, in the above situation.

While I suppose it's abstractly conceivable to test all triples of variables in the universe to see if they fit the above conditions, there is a way that won't take so long.

What we need to do is start with behavior as it has been observed under the old interpretation, which is diagrammed this way:

FIG 1.

```
       ------------                 -------------        --------
      |           |     |        |
      |  BEHAVING  |--->RESPONSE-->|ENVIRONMENTAL|----->|CONSE-  |
      |   SYSTEM   |      |        | TRANSFORMA- |      |QU¦+~?ENCES |
      |            |      |        |    TIONS    |      |        |
       ------------       |         --------------       --------
            ^       medium-low
            |       correlation
            |             |
            |             |
         STIMULUS --------
```

In fact the arrow between stimulus and behaving system is seldom verified as existing. The stimulus is defined on the basis that there is a response correlated with it. So basically we have some variable, situation, or other entity in the environment defined as a stimulus on the basis that when it occurs, the behaving system exhibits a response.

If in fact the "consequence" above is under control, then what is usually interpreted as the stimulus variable is really a disturbance, and the actual stimulus is the "consequence," as shown below:

FIG. 2

```
                           (actual sensory path)
              ----------------------<-----------------------
             |                                              |
       ------------                 -------------        --------
      |            |               |             |      |        |
      |  BEHAVING  |--->RESPONSE-->|ENVIRONMENTAL|----->| CONSE- |
      |   SYSTEM   |      |        | TRANSFORMA- |  |   | QUENCE |
      |            |      |        |    TIONS    |  |   |        |
       ------------       |         --------------  |    --------
                     medium-low                   high       |
                     correlation                 negative    |
                          |                     correlation  |
                          |                         |        |
         STIMULUS --->---------------------------------------
```

Note that for some environmental transformations, especially variable ones, the "medium-low" correlation will be very low indeed, while the "high negative" correlation will remain high. Note also that the high correlation is between EFFECTS (lines) and not variables (boxes). The stimulus (a variable) may exert a force (line) on the consequence (another variable, such as a position), while the response (really a box) acts through some environmental transformation (a nonlinear rubber lever) to produce a second force (line) acting on the consequence (position). The force exerted by the lever will be nearly equal and opposite to the force exerted by the stimulus variable: the two forces will show a very high negative correlation. The response variable will show a medium-low correlation with either the stimulus variable or its effect. By medium-low, I mean in the range normally accepted as significant, but far below 0.95.

In the above form, the relationships are clear: the stimulus is really a disturbance of a controlled variable, a consequence of the response. A more illuminating arrangement of the diagram may show what the problem is more directly;

FIG 3.

```
       ------------
      |            |
 <----|  BEHAVING  |---<-----
|     |   SYSTEM   |         |
|     |            |         |
|      ------------          |           |
|           ^           sensing of       | This part unknown to
|           |       controlled variable  | conventional
|           |                |           | behavioral-psychological
|      ------------          |           | research
|     |            |         |           |
|     | CONTROLLED |---->----            |
```

```
|---->|  VARIABLE  |                      |
|     |            |                      |
v      ------------                       |
"RESPONSE"  ^
            |
            |
       DISTURBANCE-
           OR
       "STIMULUS"
```

Now, unbeknownst to the experimenter, there is some aspect of the environment between the apparent stimulus and the behaving system. This could be something observable, or it could be inside the behaving organism. It is affected by the "response" of the system, and also by the supposed "stimulus." The real stimulus is what is sensed of the controlled variable, but the apparent stimulus is the one shown at the bottom. Not being aware of the possibility of controlled variables, the experimenter (a) assumes that the STIMULUS is directly affecting the senses of the organism, and (b) that the response is simply the result.

So the search for controlled variables can begin with existing observations of behavior and its apparent dependence on environmental variables. A more detailed examination of the experimental situation, plus suitable manipulations that interrupt the effects of the system's action on the controlled variable and/or the ability of the system to sense the state of the controlled variable, will show whether the actual situation is that of Fig. 1 or Fig. 2. If it's found that Fig. 1 applies, we have control, not reaction.

For higher-level variables, a slower time-scale has to be adopted so that action and disturbance can be seen as concurrent.

This approach does not force all behavior to be interpreted as control. It simply opens the possibililty of control and shows how to differentiate control (Fig. 1) from reaction (Fig. 2).

As to William James, CSGers (including myself) have been citing him for some years as showing clear recognition of the phenomenon of control. This only serves to show the difference between control as a phenomenon, and control theory.

Language later.   Best to all   Bill P.

Date: Wed Mar 11, 1992 8:57 am PST
Subject: Re: Low correlations

[Martin Taylor 920311 11:15]
(Bill Powers 920310 17:00)
>
>Martin Taylor (920309), Bruce Nevin (920310) --
>
>You guys are going all around the point here. I have a distinct feeling
>that you're avoiding it. On the other hand, there may be some critical fact
>that I don't seem to be communicating, so let's try again:
>

I'm not sure which way the lack of communication goes, but it sure seems to be there. I see nothing in your posting, whether by diagram or in the explanations that is not crystal clear, and has been so for ages, at least in my mind, if not in my writings. It's the starting point for most of what I have been trying to write over the last couple of months. What might be the point that Bruce and I are missing? Could you put it another way?

I think that if you take your posting together with the following paragraph from mine of 920309 11:00, you may see why I think we have a communication problem, and perhaps will be able to resolve it:

>If, as an experimenter, one can presume some pattern in the mutually observable
>environment represents a perceptual variable being controlled by the subject,
>then one can attempt to disturb that pattern and see whether the subject
>acts so that the pattern is restored or maintained. The pattern will show
>little correlation with the experimenter's disturbances or with what the
>experimenter observes of the subject's actions. If the experimenter happened
>to be correct that what she did would have disturbed the pattern if the
>subject had not been there, then there is evidence that the subject is
>controlling. The presumption that the experimenter would have disturbed the
>pattern is just that, a presumption. It is not an observation, because it
>didn't happen. Explaining why things do not happen is trickier than providing
>rationales for why they do happen. The failure of a presumed "cause" is
>easier to justify as that it was not a cause than as that an exactly
>countervailing cause was applied at the same time. I think this is at the root
>of the communication difficulty with cause-effect psychologists. Causes have
>effects, and PCT is supported when what should be causes are observed to
>have little or no effects.

Note the words in the 6th and 7th lines: "If the experimenter happened
>to be correct".

At the same time, I would love to pursue my degrees-of-freedom discussion, but I can't until the question of zeros is resolved. Could you explain the matter of the error signal for sequence, which is one part of the remaining problem? Rick has partly resolved the other part--that the spreadsheet provides a counter-example--by pointing out the considerable non-orthogonality among the ECSs in the spreadsheed. The necessary non-zeroing related to non-orthogonality was to be part of the later discussion, but I had forgotten that it could come up even when the degrees of freedom for input and output of the hierarchic CS are in balance.

Martin

Date: Wed Mar 11, 1992 11:06 am PST
Subject: Language

RM79/[From Bill Powers 920311.0930)]

Bruce Nevin (920309) --

>... it was unclear whether you were offering a game of "let's
>you and him fight" or conciliating "Boys! Boys! Don't fight!"

More the latter. The motivation, however, was to challenge you and Avery to compare the assumptions and methods on which your two approaches rest.

These assumptions and methods must be very different, assuming different models of language processes in the brain. I'm hoping for some comment on my comment to the effect that "you can't both be right and either approach could be wrong if you're trying to describe language universals." I'm hoping that you will both try to see what you're doing, in this process, as control of perceptions, and elucidate what those perceptions are. Bruce: How can you tell when you have a satisfactory expansion? Avery: How can you tell when you have a satisfactory parsing? That is, what do you look at to see whether the result meets your intentions? And what are the intentions? In short, I'm trying to get a discussion going at the next level up, rather than spinning out more examples that keep the superordinate perceptual control systems in the background. I understand that you both have theoretical cranks you can turn which will grind up a sentence and spit out an analysis according to some procedure that has been fabricated to produce that analysis. I want to get off the subject of what is spit out and get our attention onto the grinding machines. It is highly unlikely that I will be able to contribute to your efforts in linguistics by examining the outputs of these machines.

Both methods, as far as I can see, depend on some lexicon in which the characteristic uses of specific words in specific contexts are listed. It seems to me that this is a level of perception and control that can be dealt with independently of higher operations that are done once the lexicon is available. So far it seems to me that the modeler/theorist is supplying this lexicon out of informal private experience and knowledge (either you know what a verb is or you don't), instead of from a publicly-defined model. If a model satisfactory to both parties for the development of a lexicon can be sketched in, or more than sketched in, it seems to me that we would have some intermediate parts of a hierarchical model of language that would have a better chance of universality, at that level, than the greatly divergent higher-level processes that are applied using the information in the lexicon. Perhaps by making the lower levels as explicit as possible we can find reasons for whatever disagreements remain at higher levels.

Your examples of the way in which physical production of sounds influences the way phonemes are heard and used point toward a very low level part of the model that, I think, can be specified reasonably well (well enough to go on with). I'm now talking about specifying a slightly higher-level blob in which we take word production and perception for granted up to the level where the word is agreed to exist (even though it may be subject to different higher-level interpretations), and become concerned with the most elementary level of attaching words to meanings. What kinds of words get attached to what kinds of perceptions? This is not a complete lexicon, because if we take the least possible upward step we will not reach categories such as "noun" or "verb" or "operator" or "argument." That will come later.

I'm proposing that we use the same method I used in building up a systematic guess about levels of perception in general. The idea is to peel off layers, from the bottom up, that seem self-contained enough to become the units perceived and manipulated by the next higher level. Sometimes, Bruce, you refer to a back-and-forth interaction between language and meaning, providing a vague picture of some very busy multileveled process in which things are going on at many levels at once. I think we can do better: I think we can pick out those processes that occur at one level, with higher level processes OF A DIFFERENT KIND going on at the same time.

The higher-level process does not have to handle the processes going on at lower levels, only the processes that are of a new and superordinate type. Conversely, if we can find well- defined packages at lower levels, they will not have to handle aspects of language that higher levels will later be found to handle. What we will have at any given level of this kind of peeling-off process will not be language itself, but the foundations of full-blown language. And as we define the lower levels, what remains to be handled will become clearer and clearer. As we keep going up by the smallest steps we can think of, adding the least increment of function that seems to hang together, the whole structure will come to look more and more like the language we know.

It may be that a lot of the confusion in linguistics is due to trying to handle different levels of processes as if they were mixed together at one level.

>If you accomplish the aim of accounting for what all languages have in
>common, and you show that it all comes down to characteristics of the
>world of nonverbal perception plus fundamentals of physics and
>chemistry in the environment, like the acoustics of the vocal tract--
>having reached the state where linguistic universals are trivially
>deduced from first principles, what would remain?

Nothing. I think you're pulling back from reductionism, which isn't implied by my suggestion. If we find true universals, I would expect them to include such things as the capacity to recognize and execute programs in which both symbols and continuous variables are arguments and outputs, or the capacity to generalize and perceive principles. What are you doing in the search for language universals but trying to perceive principles? How do you do it, but by applying rules and algorithms at the program level? And why do you do it, but to construct a system concept of language? The hierarchy of perception and control contains what the linguist is doing at many levels, and it probably also contains language itself which is, after all, something we do with our brains.

I don't claim that the conventions of language will be "simple and uncontroversial," any more than I could claim that any other human conventions are simple and uncontroversial. We can think in either simple or complex ways, and our conventions can be easy or difficult to comprehend and agree on. But we will find it easier to agree on what the logical conventions are if we can remove lower-level aspects of language that don't depend on the program level.

Language shows us the sorts of things that a brain can do. These things are more universal than language. But the study of language gives us a window into the higher-level processes of a brain -- if only in the form of elaborate models constructed by linguists. EVERYTHING ANY HUMAN BEING DOES IS EVIDENCE FOR A MODEL OF THE BRAIN. The conventions of language tell us about the human ability to perceive and control for conventions. They tell us first that human beings in general use conventions, and second that students of human behavior can also perceive those conventions, and presumably control for conformity with them. There is no privileged position from which a linguist can see these conventions without using the very same capacity of the brain. The linguist is in no better position to grasp the conventions that others use than to grasp the conventions the linguist is using. The very perception of "convention" itself demonstrates a function of the linguist's brain.

>In particular, I believe that operator grammar shows a simple structure
>for language--a structure of word dependencies--that is universal and
>that accords well with perceptual control,...

I agree that it does, although you will have to agree that it doesn't completely fit natural language as it is spoken without introducing some important invisible processes which are in principle unverifiable. One of the things the brain can do is create plausible sets of rules that appear to fit what is observed. Often achieving a fit requires imagining information not actually present in perception. The imagined information is whatever is required to make the rule fit what is observed.

The most convincing models are those that require us to imagine the least while still fitting what we actually observe. Operator grammar requires us to imagine some critical parts of the process of language comprehension. Avery's approach requires us to imagine other kinds of hidden processes. But in either case, the rules can be made to work if we agree to imagine as prescribed.

Given any set of experiences, it is possible to devise a rule that fits them. This is like curve-fitting, only more complex. We need a way to find out whether a given "curve" has some underlying justification, or whether it is simply one of an infinity of curves that would pass through the same data points. When we compare different sets of rules for dealing with the same observables, and when neither set of rules fits the observations without adding some imaginary data, we then have to ask which rules require the least imagined data to make them work. We have to examine the imagined data to see if some of it is more believable, or if some is in principle more testable, or if some seems to be needed not just for these rules, but for others in different universes of discourse.

So I ask both you and Avery: in your models of language structure, which parts of the phenomenon of language are observed, and which are imagined in order to make the analysis work?

Avery Andrews (920310) --

> ... the pre-adapation for language is guessing what people are up do
>by watching what they are doing. Given this, one can convey intentions
>by miming, from which arises manual sign language. Then one has to get
>from that to spoken language, a step which strikes me as very
>mysterious.

By "what people are up to" I take it you mean "what people are controlling for" -- that is, what the movements they make are intended to accomplish. At the miming level, you simply take the movements as controlled variables and learn to control them for yourself. But once you've mastered the movements well enough, you have to go up a level and ask what they accomplish, what higher-level variable is controlled by varying those movements (or more generally, controlling those lower- level perceptions) that you now know how to control. So now I can say "da" and "ba" and "ma" and "baw" and "boo": that was fun, but so what? What do I use them for? Ah, you're showing me that round red thing and saying "baw." I will show you the round red thing and say "baw." Now I mime you at a higher level. If I want to see the round red thing I will say "baw" and see if that works. If I show you the round red thing you say "baw" -- or something pretty close to it. So if I want to hear "baw" I can show you the round red thing, or if I want to see the round red thing I can say "baw." If you were a different parent, say a deaf one, I wouldn't learn to say "baw" but to make a configuration with my hands. Then I could learn to use that hand configuration to get a round red thing from you, or show you the round red thing to make you do the hand configuration again. Manipulating either experience at the lower level thus becomes a means of controlling for the other. The environmental link, in both directions, consists of the relationship the parent is controlling for such that the word is produced on seeing the object, and the object is produced on hearing the word. I'm being taught, but I don't know it. I'm just learning to manipulate some things in order to control others, which is the most fun there is.

I think this is how we should build up the model for acquiring a lexicon (see above comments).

Best to all,

Bill P.

Date: Wed Mar 11, 1992 12:43 pm PST
Subject: CT details; Coin Game

[From Bill Powers 920311.1100]

Martin Taylor (920311) --

>I see nothing in your posting, whether by diagram or in the
>explanations that is not crystal clear, and has been so for ages, at
>least in my mind, if not in my writings.

>I think that if you take your posting together with the following
>paragraph from mine of 920309 11:00, you may see why I think we have a
>communication problem, and perhaps will be able to resolve it:

>[...] If the experimenter happened to be correct that what she did would
>have disturbed the pattern if the subject had not been there, then there
>is evidence that the subject is controlling.

My ears pricked up at this. Technically, you're right -- however, the condition isn't that there would have been an effect if the subject hadn't been there, but that there would have been an effect if the subject's actions hadn't canceled it. The subject might have been controlling for some other aspect of the environment. The critical thing is that you disturb something that seems to be affected by the subject's action, but the disturbance isn't counteracted.

Play the Coin Game, please. The Test is done WITH the subject there. The actions of the subject can be perceived as affecting the environment in many ways, and objectively has many different effects on objects, relationships, etc. in the environment. The question is which, if any, of these effects of the subject's actions is under control. The experimenter devises a disturbance that will alter one of those effects. If the effect changes -- if the subject does not change the action in a way that prevents the change from taking place -- then that effect of the action is not under control.

The Coin Game:

Use four coins (same or different as you please). Two people play, an Experimenter and a Subject. The Subject places the coins on a table such that they exemplify a pattern or condition that the subject has in mind. The Subject privately writes down this reference pattern on a piece of paper, and hides it. The Experimenter is to discover what the controlled pattern is, by means of disturbing the arrangement of the coins.

The rules are as follows. One round of the game starts with the Experimenter doing something that alters the arrangement of coins on the table. The Subject looks at the new arrangement, and if the target pattern can still be seen, says "No error." If the pattern now differs from the target pattern, the Subject makes any rearrangement of the coins required so that the percieved pattern once again matches the target pattern. After either a "No error" response or a corrective move, it is the Experimenter's turn again.

The game ends when the experimenter can demonstrate three different moves predicted to produce a "no error" response, and three different moves predicted to produce a correction. Then the subject displays the written description of the reference condition. No verbal communication except the words "no error" takes place during the game.

You might think at first that it will be easy for the Experimenter to discover the pattern, and compensate by choosing (as Subject) a complex reference condition. I advise choosing a simple reference condition if you want the game to finish in under half an hour, or not be abandoned.

This game illustrates all the facets of the test for the controlled variable. Clark McPhail has been using it to teach The Test (he sent me copies of the experimental reports by about 50 of his students -- wonderful reading, especially the comment by one student that he really admired sociologists for being able to use The Test in their work, because it is so complex).

Martin, I know you have a deep grasp of control principles. I expect no less of you. But however much one knows, there are always blind spots and misinterpretations. Control theory reveals endless depths of new detail, and I don't know of anyone who has plumbed them completely in just a few years. Ask some of our Old Hands when the last new understanding came to them, and how long they'd been control theorists before that. And think of how much remains to be developed, that nobody has answers for!

Zeros and sequences:

Suppose you want to produce a sequence like "now is the time for all good men to come to the aid of their country." By the wasteful pandemonium postulate, this is the province of just one sequence-recognizer and control system.

I want to produce this effect: the sequence is recognized while it is occurring, maximum recognition resulting when the sequence proceeds exactly as the system is designed to recognize it, to the end. The reference signal from a higher system doesn't contain this sequence (if it did, the higher system would be the sequence level!), but just says "make your sequence appear in perception." The perceptual signal isn't the sequence, but simply an indication of the degree to which the sequence in question is occurring. The sequence is physically present in the sequence of reference signals and subsequent perceptual signals in individual lower-level systems that supply the elements. So we recognize the sequence, and can correct errors when they occur, and can produce the outputs that keep the sequence going. It's hard to imagine how to produce the correct sequence of outputs without defining the sequence twice: once in the recognizer, and again in the output function.

Clearly, this system needs to be able to generate a perceptual signal indicating that some particular sequence is in progress so far. It also needs to be able to switch lower-level reference signals to produce the elements of that sequence, which are lower-level perceptual signals. Disturbances of the sequence (incorrect elements) should produce an error and a change in the output that does something appropriate -- resets the whole system, replaces the wrong element, and so on. I don't know how to design a system that will behave exactly like this. I think it could be done, but doing it with the sequence defined in only one place would be tricky. The design will probably not look like three boxes.

This doesn't really answer your questions about zeroing, but it does answer one question: does Bill have a design for a sequence-controlling (as opposed to "emitting") system? The answer is no. There's some principle missing here. These higher-level systems are hard and fuzzy at the same time.

Best   Bill P.

Date: Wed Mar 11, 1992 4:27 pm PST
Subject: Re: CT details; Coin Game

[Martin Taylor 920311 17:30]
(Bill Powers 920311.1100)

Yes, I could have been more precise: the subject being there is an insufficient condition. But the point of the argument is that control cannot be detected in the absence of a prediction about what would have been disturbed in the absence of control.

I appreciate the Coin Game. My thesis supervisor did quite extensive studies of this kind of game, perhaps not structured exactly the same way, but very like. He was studying perception, not control, but the issue is also of determining by such trials the nature of prespecified relationships. I think it was a popular experimental paradigm at the time. The effects of interactions on the kinds of relations that were readily detected or were detected only with difficulty was the point at issue. He came up with the notion of integral and separable perceptual dimensions, which seem to be quite important. I can't give any specific references, but if you want to search for them, look for W. R. Garner in the late 50's or 60's.

>Martin, I know you have a deep grasp of control principles. I expect no
>less of you. But however much one knows, there are always blind spots and
>misinterpretations. Control theory reveals endless depths of new
>detail, and I don't know of anyone who has plumbed them completely in just

>a few years. Ask some of our Old Hands when the last new understanding came
>to them, and how long they'd been control theorists before that. And think
>of how much remains to be developed, that nobody has answers for!

I think you credit me with more than I deserve. There are indeed many puzzling holes, and most of my postings are either trying to expose them for you to fill, or to fill holes I perceive in the understanding of others. At the moment, the hole that dominates my puzzlement is the one you began to address in this posting, of how to assert error when the relevant percepts are extended in time. I'm not clear what it means to have a sequence error, though it is clear what it means to have a sequence that is in error.

In this kind of slow interaction (fast by comparison with that of Abbe Mersenne, but slow compared to sitting around a blackboard), the problem is to determine just what is the problem. It is often clear that a misunderstanding exists, but the attempt to clarify it often misses the mark, and sounds like "teaching your grandmother to suck eggs." I don't know of any solution to this problem, except to rely on the perception that any such attempt (a) was well intentioned, and probably helps other readers besides, and (b) can point the way to a better understanding of where misunderstandings really lie. We have to put up with it.

Martin

PS I have just been informed that our system will be cut off from e-mail between mid-afternoon March 12 until some time March 16. So don't anticipate much from me in that period!

PPS Did anyone (Gary?) ever contact Jan-Olof Eklundh?

Date: Thu Mar 12, 1992 9:52 am PST
Subject: Re: motor programs, caused output, the curse

[Rick Marken (920312)]

I said (Rick Marken 920310) in moderate jest:

>Motor programs? Responses determined by input probability?
>Gary, help, the curse of PCT has taken over CSGNet.
>Anybody want to talk controlled variables, maybe????

And Martin Taylor (920310) replied:

>I don't think you can simply throw away all of perceptual psychophysics
>because
>the subject does not control what the experimenter presents. The controlled
>percept is at a higher level, and the experimenter knows it.

Well, I don't know that the experimenter really looks at it that way. I used to do auditory psychophysics and I looked at the whole thing in s-r terms. Yes, we did think about the subject's response goals (I used SDT, of course). But even then we tried to see what was probably equivalent to a reference signal in SDT (the criterion) as a response to input -- the payoff matrix. But I agree that we did understand that the subject was controlling (probably) at least one variable -- response probability or (more likely) sequential response proabbility. I guess we also know that, under "ordinary circumstances" (no monetary payoffs) the criterion setting was secularly adjustable (just like a PCT reference signal).

I didn't mean to imply that this work was valueless. I was being a bit sarcastic. But there was an element of seriousness there. I do think that it is probably a waste of time to sift through the misconceptions and fairly useless data of conventional psychology. I think it's just a waste of time looking for areas of convergence between PCT and conventional psychological research. . I think fair is fair -- conventional psychology is pretty content to ignore PCT; I think PCT should just return the favor. I think we have already shown the fundemental flaws in the conventional approach and given precise, alternative explanations of what conventional psychology considers to be some of its major phenomena (operant conditioning, coordinated behavior, reflexes, etc). My personal feeling is that PCT people can be much more productive by just starting from scratch, pretty much, rather than trying to deal with a world of observations and theories that were based on the wrong assumption about the nature of behavior. But that's my personal opinion -- those of you who want to try to apply PCT as an alternative to "conventional" models are free to try it; but as you begin to get a deeper and deeper understanding of PCT I bet that you'll give it up and just start studying control phenomena directly.

>Does that take the curse off, a teensy bit?

Of course. Pay no attention to that loose canon behind the keyboard.

Hasta Luego   Rick

Date: Thu Mar 12, 1992 12:49 pm PST
Subject: progress report on arm model

[From Bill Powers (920312.1000)]

For a couple of weeks I've been back on the Arm program Version 2 (in C), trying to get it in shape to go out the door. It's now working about as well as possible -- not perfect, but usable and instructive. The next step is to send it on to Greg Williams, coauthor, who has some tests he can apply to it. Greg and I will try to get a paper together which we will send to Science (my letter on an open-loop model by Bizzi was rejected, mostly because the referees (!) objected to my references to an unpublished arm model ... the implication being that I should put my model where my mouth is. So I will.). This version does not yet have nonlinear muscle spring constants, nor does it use the available data showing how muscle force converts to torque at various joint-angles. I've tried those features and they seem to work, but I'm burned out for now. They can wait for Version 3.

There was a big breakthrough this week. To explain it, I have to describe the basic spinal-cord part of the model. I won't go into every detail.

The kinesthetic control systems in this model are taken directly from the basic known circuitry of the tendon and stretch reflexes. Schematically:

```
                              Alpha efferent
  Gamma efferent             (voluntary reference
 (length reference               signal)
```

```
     signal)                      |
       |                          |      (Spinal
       |        (stretch          |       Motor
       V        error sig)      + |       Neuron)
   Mechanical                 +   V     /
   Comparator --- Kg ---------->Comparator------------
       ^          ^        ----> -                    |
       |          |       | (force signal)            |  (contraction
       |          |       |                           |  sensitivity)
       |    (sensitivity)->--Kt--      (spring       Ko
       |                         |      constant)     |
       |                         |force    Ks         V MUSCLE
       |             anchor |||--O------- /////////////////--LOAD
       |                        receptor    --||||||||
       |                                   |   stretch
       |      (mechanical effects)         |   receptor
        ---------------<-------------------
```

This diagram really represents two control systems using opposing muscles and balanced (push-pull) pairs of reference signals. I have reduced the pair to a single system to make the model simpler. As a result, the "muscle" can both push and pull on the load, and all signals can be both positive and negative. Load position is an angular measure.

There are two loops in this system: the force loop involving the Golgi tendon organs and the stretch loop involving the muscle spindle containing the annulospiral stretch receptors.

The lower loop is a force-control system. The net reference signal is the sum of the stretch error signal and the voluntary reference signal, the alpha efferent signal. This system maintains the sensed force at the net reference level, with a loop gain of KoKsKt (when the load point is fixed). Pulling on the load point will cause the muscle to shift the load point in the direction of the pull (one muscle will pull in the direction of the load while the other relaxes). As a result, the sensed force in the tendon will remain the same. This system reduces the effective mass of the arm nearly to zero. It will apply the specified force regardless of the position of the load.

The stretch loop controls the length of the muscle and thus the load position. The feedback signal is a rate-plus-proportional signal, with the two components independently adjustable (not shown). The mechanical comparator is actually in the muscle spindle, the gamma efferent signal shortening two muscles that stretch the annulospiral receptor supported between them. Thus either an increase in the gamma efferent signal or a stretching of the main muscle will cause a rise in the stretch feedback signal. This feedback signal is actually the error signal.

The higher or outer control loop causes the muscle length to follow the gamma efferent signal, by means of altering the reference level for muscle force. Muscle force is actually equivalent to angular acceleration of the arm. So the outer loop controls angular position of the arm by adjusting angular acceleration, with the rate component of position feedback damping the system.

That's the basic idea. The remainder of the model is a visual control system using binocular vision for depth. A target and a fingertip are imaged on the retina, with x and y positions of the images being computed by an abracadabra box. The outputs of the visual control systems go to three spinal control systems as above: the depth error varies the elbow angle, the y position error varies the vertical shoulder angle, and the x position error varies the horizontal shoulder angle. There is no twisting degree of freedom at the shoulder. The visual control systems use proportional-plus-integral control.

Until last week, I used the gamma inputs as the reference inputs to the spinal systems. After all, the gamma input is a position reference signal, and on casual inspection I decided that either of the reference inputs would have the same effect. So I just set the alpha reference signal to zero. The model worked, but it seems skittish, hard to stabilize. It couldn't move very fast without driving some signal to a limit (I used fixed-point arithmetic in the visual systems). I kept fiddling with it because it wasn't pretty.

Then last week it occurred to me to try using the alpha reference signals instead, just to see what would happen, fixing the gamma reference signals at constant values. Wow. Suddenly I could get the arm to travel through a full-scale movement in 0.1 second and it was an order of magnitude more stable than before. Then I figured out the problem. The gamma input path contains fast rate-of-change feedback for stabilization. When I hit the gamma reference input with a square wave (the testing mode), the rate part of the stretch signal went sky-high and caused all sorts of problems. Changing the input to the alpha input couldn't cause the muscle length to change that fast, because of the mass of the arm, so the rate signals behaved themselves.

I should have tried that right away. The alpha reference signal is the input from the voluntary systems in the brain; the gamma reference signal is not. Now I realize that the gamma input is just for slower reflexive adjustments (as I have read many times), such as setting a base arm position and holding it against gravity while the alpha system varies the arm position around the base. The alpha reference signal initially changes the force being controlled, accelerating the arm. But that changes the muscle lengths, and the gamma system alters its reference input to the force control system, bringing the arm to a stop in a new position. It's perfectly clear now, but I just didn't see it. I was fooled because the alpha input would seem to control force or acceleration, not position. But because of the gamma system, it ends up altering position in a stable and controlled way!

I made one more change to the model. When the target jumps to a new position, a "virtual target" is computed which moves to the new position at a constant fast speed (adjustable) in x, y, and z. This virtual target is the one that the visual system follows. This is equivalent to the operation of a higher-level system that perceives in objective coordinates. The moving virtual target is like a reference signal for position in objective space that moves toward the new target position, dragging the finger along behind it. The result is that the finger is always moving toward a reference position in objective space that is on a straight line connecting old and new target positions.

Before doing it this way, I put some perceptual computations into the model to compensate for the visual distortion of an xyz space into an r-theta-phi space. This worked, more or less, but clumsily. As I didn't want to get into a real design of the fourth level of this model, I decided to set up a system equivalent to what the real perceptual system would accomplish, which would be to perceive the fingertip as following a reference position moving in a straight line toward the new target position.

Now the motion of the fingertip toward the new target position is controlled at all times, with only a modest amount of error even for large movements taking only 0.25 second. For slow motions (as when the fingertip describes a circle around the target), the errors are very small. It draws a nice round circle.

The screen now shows the joint angles and the torques applied to each joint. When the target makes a very large jump, the relevant torque rises until the arm velocity is keeping up with the virtual reference position, then falls nearly to zero. After a short "ballistic" interval, the same torque peaks in the other direction, slowing the arm to a stop. So now the arm shows the kind of behavior that has been noted qualitatively in the literature -- but control is continuous and there is really never any ballistic motion. The motion seems to be ballistic only because the arm is coasting at nearly the same speed that the reference position is changing and there is little error to produce any muscle force.

Another rather astonishing effect showed up when I started using the right reference input for the kinesthetic systems. If the visual feedback systems are cut off from the kinesthetic systems (in a testing mode), and if gravity is turned on, the arm very slowly sinks toward the straight-down position. This is the condition known as "waxy flexibility"; it results from loss of the reference signals from supra-spinal systems, as in a spinal transection above the level of the control systems. The cause is the interplay between the gamma loop and the alpha loop -- if you look at the diagram, you'll see that the feedback signs from these loops are opposed. The best values for the constants Kg and Kt, the stretch and tendon sensitivities respectively, are largish numbers (30 to 60), equal to each other. The equations show that when Kg and Kt are nearly equal, there is very little active position control! I thought at first that I'd inadvertently broken some control loop in switching to the test mode -- but then I remembered something I'd read years ago about types of paralysis, and the descriptive name, waxy flexibility, popped up (I had remembered "waxy immobility" but Mary remembered the right name).

Date: Thu Mar 12, 1992 6:07 pm PST
Subject: Genetic algorithm in behavioral modeling

[From Bill Powers (920312)]

Hello, Randy --

Thanks for the paper on the Genetic Algorithm and your use of it with dynamical neural networks. As I understand it, you map selected characteristics of a neural network (like connection weights) into a "genome" in which each "gene" is a small number of bits. Then by using a model involving random mutations and cross-exchanges of genes between reproducing pairs, you get variations in the behavioral characteristics from generation to generation. There has to be some selection pressure to weed out incompetent individuals -- as I understand it, in your model this criterion would be whether a given individual can get to a food patch by some sort of chemotaxis (I'll limit my remarks to the first model -- my comments would apply as well to the locomotion model).

The model of chemotaxis you start with (I'm repeating all this both to check my understanding and to summarize for others on CSGnet) has a body with two chemosensors on it and two effectors one of which moves the body while the other turns it, with velocity proportional to force. There's a six-node neural net inside, receiving the two chemical signals and emitting the two motor signals. The connectivity of the network allows 24 parameters to be adjusted, with 4-bit accuracy; a 96-bit "genome".

To evaluate fitness, you do many runs with varying starting conditions, and record the squared distance from the food patch at the end of a set number of runs (puberty?). While you don't describe what happens next, I presume that if a fitness threshold is passed, "reproduction" takes place among the survivors, with mutation and gene-swapping to produce individuals for the next generation of trials. The ones that don't reach the survival threshold are deleted.

I salute the cleverness of this approach; clearly, given enough trials and a small enough genome, it will autonomously produce individuals capable of meeting the criterion. I also want to point out that this model is probably related to evolution more in the manner of an allegory than of a description, which you probably realize. It really doesn't make literal sense to select for anything but successful arrival at the food patch -- merely being "closest" can't have any effect on natural selection. It's the experimenter, not simulated nature, who decides to reward those who got closest to the food by letting them reproduce. I suspect that the reason for these rather lenient definitions of fitness is that using a realistic definition (get to the food or die) resulted in immediate extinction of all populations. Without some external intelligence watching progress toward the required behavior and meting out reward and punishment according to the "best tries," the Genetic Algorithm would never arrive at the correct result save through chance mutation (at least in this application). And even then, arrival may have been accidental and unrepeatable even using that same organization.

This raises a real problem for evolutionary theory and its application to complex organizations, which I'm sure I'm not the first to point out. Whatever the adaptation, it must actually go all the way to success if the real fitness criterion -- living to reproductive age -- is to be met. The bug that succeeds in traversing only 99 millimeters of the 100 millimeter distance between it and food will fail just as surely as the one that starts immediately by going the wrong way. The fitness criterion of natural selection is so crude that it can't distinguish degrees of success short of complete success. If your chemotactic system is taken as a test of natural selection as an evolutionary theory of behavior, I think you have shown that this concept of the origins of complex behavioral organization is weak. It can't work without an external observer who can see which directions of change would be beneficial if given a chance to go further in that direction.

This does not mean that the basic approach of GA is impractical -- I think it is just mis-named. Everything that the GA approach does by applying fitness criteria to successive generations of organisms can be done by using a slightly different approach within a single organism in one lifetime, an approach that I call "reorganization." I don't mean that reorganization can substitute for natural selection. But it can select for things like complex behavioral organization that strict natural selection hasn't a hope of achieving. What is achievable through natural selection, I think, is something more basic: production of organisms that can reorganize. Until that level is reached, evolution will be extremely slow and organized behavior will be very simple, much simpler than your chemotactic organism.

Reorganization, I now see, works a great deal like the genetic algorithm, but without the genes and without a binary pass/fail criterion. The "fitness criterion" is a set of "intrinsic variables" being compared with a set of "intrinsic reference signals." These variables can be defined much as you would define fitness criteria, except that they are continuous variables and not just binary events. The "intrinsic reference levels" are the states of those variables specified as the "best" state. In fact, these variables and their reference levels can be considered to be products of natural selection in the usual way.

The result of comparing intrinsic variables against reference states is a composite signal I call the "intrinsic error signal." This error signal drives the process of reorganization. I speak here as if there is a single reorganizing system, but the same principles would apply if there were multiple reorganizing systems associated with smaller subdivisions of the organism -- that's a matter best left open for research and creative modeling.

The reorganizing output is a strictly random change in whatever part of the system is affected by it (changes, for example, in parameters like the ones you link to the bit-string "genome"). But it is not random in one and only one respect: the rate at which the random changes are caused. That rate is proportional to the intrinsic error signal's absolute value.

The overall effect is that when intrinsic error is large, random changes of organization are caused frequently. If one of those changes results by any means (or even by accident) in a reduction of the magnitude of intrinsic error, the interval between reorganizations lengthens. Thus the organization that was produced last persists a little longer before the next reorganization. If there are systematic relationships between the parameters being randomly varied and the various intrinsic variables, then the result will be a biased random walk that will end with a state of very small intrinsic error and an organization that is only infrequently changed (assuming the organism doesn't die first).

This principle of reorganization was described in my 1973 book, Behavior: the control of perception. It was an elaboration on W. Ross Ashby's principle of "superstability" and Don Campbell's prior notion of "blind variation and selective retention." But it wasn't until the 80s that I found evidence that this kind of system would actually work efficiently enough to be considered as a realistic possibility. Oddly enough, what convinced me was a book on chemotaxis by Daniel Koshland.

Koshland studied E. coli, which exhibits efficient chemotaxis without being able to detect the spatial direction of chemical gradients or steer. All E. coli can do is swim at a constant speed in a straight line, or tumble randomly in 3-space by reversing some of its flagellae. E. coli is sensitive to the time-rate of change of certain chemical concentrations (over 20 kinds) in its vicinity. For an attractant, the rule is that when the time rate of change is positive, the random tumbles are spaced far apart, while for a negative rate, the tumbles occur close together in time (with proportionality between rate of change and delay period). In a gradient, the sensed time rate of change of concentration is determined by the direction of swimming. This is all that is required: the random walk of swimming segments brings E. coli up the gradient with about 70 percent of the efficiency it would have if it could turn up the gradient and swim that way (I did some modeling to verify this).

E. coli's behavior is an example of a "reorganizing" control system, although it isn't changing its organization. It can be modeled as an ordinary control system sensing rate of change of concentration, comparing the sensed rate with a reference specification, and producing an error signal based on the difference. All that is unusual is that the output is a random effect varying only in the frequency of its application. This kind of control system, which is essentially what I had visualized in 1973, turns out to be several orders of magnitude more effective than I had imagined when I wrote that book. It is a simple and powerful method of adaptation that works with far more efficiency and refinement than the underlying process of natural selection -- from which, presumably, this kind of system arose.

In my model, there is an acquired hierarchy of control systems, basically a neural net with many levels, organized to control the relationships between an organism and its environment. This hierarchy arises partly from inherited (evolved) prior organization, but mostly through the action of an inherited reorganizing system that is operating from the beginning of each organism's life. The reorganizing system senses and controls selected variables that indicate the status of the organism, independently of what the acquired neural organization senses and controls. The outcome of this reorganizing process is, in fact, the adult organization of the nervous system and its behavior.

The great advantage of the reorganizing process over natural selection is that the "fitness criterion" is a continuous variable, and can be multidimensional, instead of being a single crude determination of survival or failure. It does for the organism, in fact, what you have found necessary to do for your evolving model organisms: it introduces a continuous scale that tends to preserve improvements and eliminate worse behavior without requiring organisms to go all the way to absolute success.

I notice that you mention the use of Gray codes as a way of achieving small changes from bit-mutations. I have also arrived at this conclusion; that for reorganization to work, small changes must have small effects. I'm dubious, however, about the Gray code solution. True, in the Gray code, all changes in value of one least signficant digit are brought about by single-bit changes in the binary code. But the converse doesn't follow: that if you randomly change a bit, you will get only a small step-change in the value. For that minimum change to occur, you have to change the RIGHT bit, the next one that will cause the smallest transition of value. If you change any other bit, as a random change is most likely to do, you no longer get a small effect. You can actually get a change of value of any size.

Most of us on CSGnet doing modeling use analog systems. To model reorganization in an analog model, it won't do to just change parameters using a random-number generator. What we have done instead is to associate a small delta with each parameter, and let the reorganizing system's output select a value of delta at random between small positive and negative limits. Then that delta is added to the value of the parameter, assuring that any one change will be small. This means that to get a large change in a parameter, it's necessary for many successive reorganizations to produce the same sign of the associated delta. Tom Bourbon, who is on this net, has used this method to produce a self-adapting control system, and also to make a model of a control system reorganizing itself to behave like a real subject. The end-point of the reorganizing process is determined by what you define as intrinsic variables and their reference levels. For the self-adapting system, the intrinsic variable was simply the average error signal in a control system, with an intrinsic reference level of zero. In the case of matching a model to a subject's behavior, the criterion was the average difference between the model and the person on some behavioral variable, again with an intrinsic reference level of zero. Maybe Tom will (at last) favor us with a brief report on his results -- I hope I have represented his method accurately.

I think that the principle of reorganization fits between behavior and natural selection. In fact there may be several layers of reorganizing processes. This concept, I think, bridges the gap between the genes and the organization of behavior in a way that's far more believable than trying to do it in one huge jump. What do you think?

As usual, I'm sending you a copy of this direct and also posting it on CSGnet. You can reply just be sending to CSG-L as before, and I'll get it. I have a new address -- my very own logon at last -- for direct communications: see header.

Best,

Bill Powers

Date: Fri Mar 13, 1992 9:36 am PST
Subject: progress report on arm model

[From Rick Marken (920313)]

Bill Powers (920312.1000) says:

>Until last week, I used the gamma inputs as the reference inputs to the
>spinal systems. After all, the gamma input is a position reference signal,

>Then last week it occurred to me to try using the alpha reference signals
>instead, just to see what would happen, fixing the gamma reference signals
>at constant values. Wow. Suddenly I could get the arm to travel through a
>full-scale movement in 0.1 second and it was an order of magnitude more
>stable than before.

>Now I realize that the gamma input is just for slower reflexive
>adjustments (as I have read many times), such as setting a base arm
>position and holding it against gravity while the alpha system varies the
>arm position around the base. The alpha reference signal initially changes

>the force being controlled, accelerating the arm. But that changes the
>muscle lengths, and the gamma system alters its reference input to the
force control system, bringing the arm to a stop in a new position. It's
>perfectly clear now, but I just didn't see it. I was fooled because the
>alpha input would seem to control force or acceleration, not position. But
>because of the gamma system, it ends up altering position in a stable and
>controlled way!

Bill, could you help a poor old non-physicist here. It sounds great but I don't quite get it. I guess I could set up a computer simulation of the model but first let me explain my problem and maybe you can clear it up real fast: isn't load (arm, I presume) position proportional to muscle length and isn't the stretch receptor signal also proportional to muscle length? So if the gamma efferent is fixed at some value won't it keep stretch signal (and, hence, load position) constant despite variations in alpha efferents changing the requested force exerted on the load? I don't understand how variation of the alpha efferent moves the load (as you say it does in your new scheme). Wouldn't the stretch error offset the alpha efferent to keep the force at the level needed to preserve the stretch input at its gamma efferent reference? I believe it works -- I'd just like my mental model of it to work too.

Also, great post on the misnamed GAs. I'd love to hear Tom Bourbon's description of his adaptive control system. I hope this encourages another discussion of evolutionary mechanisms.

Best regards   Rick

Date: Fri Mar 13, 1992 11:14 am PST
Subject: Re: Genetic algorithm in behavioral modeling

[Martin Taylor 930313 11:00]   (Bill Powers 920312)

(One of our system administrators found a way to (manually) redistribute my mail to this machine while the one I usually use is down over the weekend, so I am not as cut off as I thought I would be. But delivery is not guaranteed, especially on Saturday and Sunday.)

Bill, I think you sell Genetic Algorithms (GAs) short. They have, over the last 20 years or so, proved remarkably robust and able to find solutions to difficult problems of many degrees of freedom with considerable efficiency. John Holland has been in much the same situation as you--an early insight that very few people noted, dogged work for a long time, and a late flowering of interest and further development in a wider circle of researchers.

The essential core of GAs is usually not mutation, but crossover. Consider an "organism" as representing an attempted solution to the problem at hand. This solution consists of N degrees of freedom (genes) that may take on k values (k may differ across genes, but we usually don't consider that; it is often set to two). In a simple crossover, two organisms mate to produce a child that has the first P genes of one parent and the final N-P genes of the other. It is important that the genes be considered as an ordered set, because minor disordering during crossover is one of the reasons why GA is a powerful soution method; genes that form a successful "team" will tend to be located near each other in successful organisms, and their nearness must both be allowed to occur by chance (reorganization) and to maintain itself across generations. I recommend reading Dawkins (e.g. Blind Watchmaker) on the importance of this in natural evolution.

More complex crossover rules are often more efficient than the simple-minded one I just described. For example, taking two cuts in the chromosome (gene string) and giving the offspring the first P and the last Q genes from one parent and the middle N-(P+Q) genes from the other is often beneficial.

Mutation likewise should be very rare in a GA. One in a thousand is not an uncommon rate to find in GA experiments.

You unfairly criticize Beer's fitness criterion, the square of the distance to food, on the grounds that the bug will die if it does not find food. This assertion depends on an assumption that the bug gets exactly one oppertunity to feed in its lifetime, and if it misses, it dies without offspring. No natural organism lives under such severe restrictions. If you miss lunch, you go hungry until dinner, but you don't die of the hunger. So, if a bug can get within a radius R of food, it is somewhere in a circle of area proposrtional to R^2, and has a chance 1/R^2 of hitting the food. Beer's fitness measure seems to me to be a measure of how often the bug actually gets a chance to eat, and to be very suitable as a fitness measure for the GA.

Also, you ignore the fact that any COMPLEX organism is the result of a long sequence of evolutionary changes, almost all of the recent ones involving little or no alteration in the low-level functions. Our basic chemistry is essentially the same as that of our cousins, the bacteria and plants. At a higher level, almost all mobile creatures are more or less bilaterally symmetric with a single central food tract and paired appendages that serve for grabbing and moving. You side with the creationists when you talk about "trying to do it in one huge jump." Evolution doesn't work that way in nature or in a GA.

All the same, GA's have successfully solved problems in which the good solution is a sharp spike in a wide landscape of non-solutions, so "doing it in one huge jump" is not always impossible.

There is indeed a lot in common between reorganization and GAs, but I think it is in the recombination algorithms of the GA rather than in the crossover and selection mechanisms. But maybe they are closer than I see.

Martin

9203C CSGnet

Date: Sun Mar 15, 1992 9:35 pm PST
Subject: Cut off

[From Rick Marken (910315)]

I have just been informed by the system manager here that all mail sent between 9:00pm friday and now (about 9:pm sunday) is in the bit bucket. Since there was some stuff I was hoping to hear about, if any CSGNet mail was posted within the above interval could those who posted it send it to me personally -- I have received nothing from the net since Friday but maybe there was nothing. If so, please ignore this post.

Thanks   Rick

Date: Mon Mar 16, 1992 9:22 am PST
Subject: Insect Locomotion

Bill Powers writes:

>> ... the artificial insect's physics isn't TOO odd. It corresponds to
>> that for a highly damped system, which probably isn't too bad an
>> approximation for legged systems)

>Well, you have to admit it's a LITTLE odd, in that the whole bug is
>damped rather than the individual legs!

But that's just my point! Muscle has damping characteristics. Insects almost always have several legs on the ground at once, and experiments have shown that certain legs generate forces which actually OPPOSE the forward motion of the body (in the American cockroach, the front legs have this property). Thus one could consider a walking insect as always being connected to the ground through one or more dashpots, which WOULD damp the motion of the whole bug.

>Greg Williams, thanks very much for the articles by Cruse . . .

I noticed the similarity between the walking stick work and your model too and I'm glad that Greg sent you some of Holk's papers. Recently, we successfully implemented his leg coordination rules for stick insects (see "What mechanisms coordinate leg movement in walking arthropods?", Trends in Neurosciences 13:15-21 for a recent review) in our hexapod robot, robustly generating a range of gaits similar to our original model.

I continue to encourage your interest in models of insect locomotion. However, it seems that you prefer to make everything purely sensory-driven. This is, of course, a perfectly legitimate approach, but I must repeat that it does not appear to be the way biology does it. A number of experiments have demonstrated that the neural circuits underlying many rhythmic behaviors (e.g. walking, swimming, chewing, breathing) can generate the basic oscillatory pattern IN THE COMPLETE ABSENCE OF SENSORY FEEDBACK. Of course, this central rhythm must be reinforced and fine-tuned by sensory feedback in order to exhibit completely normal output patterns. In addition, work on cockroach locomotion by Sasha Zill has suggested that even when sensory feedback is intact, it may come in too slowly to play any role in fast walking insects (the cockroach is capable of stepping frequencies in excess of 24 Hz!).

Best Regards,
Randy Beer

Date: Mon Mar 16, 1992 9:22 am PST
Subject: Philosophy

Rick Marken writes:

>What are "internal dynamics"? If, by this, you mean "a model of the
>nervous system" then I'd say that is certainly a way to describe the
>problem as I see it as well . . .

By "internal dynamics" I simply mean the time-dependent input/output properties of the animal. If you are interested in understanding the mechanisms underlying natural animal behavior, then this might involve modeling the nervous system. However, it needn't necessarily do so. If you are designing an artificial agent, the necessary internal dynamics could be implemented with ping-pong balls and peanut butter as long as the external behavior of the agent were appropriately coupled to its environment.

>Would the acceleration of the bug as it falls off a ledge count as a
>behavior to be modeled? If not, why not. If so, why so?

It certainly could. Some species of moths are preyed upon by bats who, as we all know, navigate by echolocation. These moths have evolved an interesting escape mechanism. Whenever they detect vibrations of a certain frequency (namely, that used by bats searching for prey), they simply fold up their wings and drop like a stone. This certainly counts as a behavior in the ethological sense, though I couldn't say whether it is a behavior in the specialized technical sense of PCT. By the way, what could possibly constitute the controlled variable in this case?

Both Rick Marken and Bill Powers summarized the basic tenets of PCT and HCT for me (which I will refer to as simply CT for simplicity). I will try to respond to their summary below. If some of my comments are based upon a misunderstanding of CT, I trust that one or both of them will correct me.

The basic idea of CT seems to be that behavior is the consequence of negative feedback control of selected sensory inputs. The use of negative feedback control to regulate important variables is clearly ubiquitous in biological systems. If this is the only point of CT, then I can't really find any reason to disagree.

However, if CT is making the much stronger claim that negative feedback control is universal and ALL behavior can be understood in its terms, then I am extremely skeptical. Animals certainly DO control some sensory inputs using negative feedback, but I don't believe that they JUST control sensory inputs. Many consequences that are of utmost importance to an animal are not controlled in any negative feedback way.

To take just one example, the American cockroach has an escape response: Whenever anything lunges toward it (i.e. a striking toad or a foot), it turns roughly 180 degrees away from the direction of the attack and runs away. A fair amount is known about the neural circuity underlying this response (we have actually been involved in modeling this system for a couple of years) and people have found that it is organized not as a negative feedback system, but as a feedforward system. Upon reflection, the reason for this is rather obvious. The sole purpose of the escape response is to make sure that the cockroach isn't where it was when the attack began. Consequently, the cockroach can complete its initial turn in about 60 MSEC. Given the latencies involved in the sensory organs and neural signal transmission, there simply isn't time to do negative feedback control of this turn. Nor is such precision necessary.

But wait, the story gets even more interesting. It turns out that the cockroach can factor a great deal of contextual information into its escape, including auditory, tactile, visual, and proprioceptive cues. Again, the necessity of this is obvious upon reflection. An attack can come at any time: when the insect is in midstride, when it is feeding, when it is near the edge of a wall, etc. The actual movements required to escape might be very different in each of these cases. For example, experiments have shown that cockroaches that are attacked near a wall will make entirely different movements than free-ranging cockroaches. It turns out that there is a population of about 100 interneurons whose job it appears to be to integrate all the appropriate contextual information and essentially always be ready to generate an appropriate escape should the insect be attacked in the next instant.

And some important variables aren't likely to be controlled at all, even in a feedforward way. Survivability is undoubtedly one of the crucial variables for any animal. But do you seriously believe that an animal explicitly estimates its current survivability and that some high level control system actually uses the error between this estimate and a "reference" level to guide its behavior? Such variables are simply too complex to explicitly estimate, nor is it necessary to do so. At best, animals may control variables that are sufficiently correlated with survivability that, on average, they do in fact survive for an extended period of time. This also suggests that many variables that appear to be explicitly controlled on casual inspection may, in fact, not be.

What I was trying to say in the summary of my own position was that evolution selects for animals that always generate the appropriate consequences PERIOD. Sometimes those consequences may be generated by negative feedback control. Sometimes those consequences may be generated by feedforward circuits (another example of this would be central pattern generators). Sometimes those consequences may be generated by the laws of physics (e.g. the plummeting moths mentioned above). Evolution selects only for the viability of the complete package and for that reason I think that it can be grossly misleading to impose our organizational preconceptions on evolved control systems.

Best regards,
Randy Beer

Date: Mon Mar 16, 1992 9:23 am PST
Subject: Genetic Algorithms

Bill Powers writes:

>Thanks for the paper on the Genetic Alogorithm and your use of it with
>dynamical neural networks. As I understand it ...

[summary deleted]

>to evaluate fitness, you do many runs with varying starting conditions, and
>record the squared distance from the food patch at the end of a set number
>of runs (puberty?). While you don't describe what happens next, I presume
>that if a fitness threshold is passed, "reproduction" takes place among
>the survivors, with mutation and gene-swapping to produce individuals for the
>next generation of trials. The ones that don't reach the survival threshold
>are deleted.

Your summary of the basic approach is accurate. However, there is no threshold for reproduction in genetic algorithms. Rather, individuals are selected for reproduction with a probability proportional to their fitness. So there is a very small chance that even a very unfit individual will get to reproduce, though the vast majority of selected individuals will be very fit.

>...this model is probably related to evolution more in the manner of an
>allegory than of a description, which you probably realize.

Yes. Genetic algorithms are about as related to evolution as the simplified neurons used in most neural network models are to real nerve cells. However, as long as one never forgets the simplifications that have been made, such simplified models can be quite useful. I would think that you would agree with this statement since you often refer to simple neural network models in your own theorizing.

>It really doesn't make literal sense to select for anything but
>successful arrival at the food patch -- merely being "closest" can't
>have any effect on natural selection. It's the experimenter, not
>simulated nature, who decides to reward those who got closest to the
>food by letting them reproduce.

This is precisely the distinction between intrinsic and extrinsic fitness that was discussed at the end of the paper. Extrinsic fitness is when the experimenter imposes an external fitness criterion of his or her choosing. Intrinsic fitness is when the fitness of an agent is directly tied to its ability to find a mate and reproduce within the simulated environment.

As mentioned in my paper, there has been some very interesting preliminary work on intrinsic fitness. For example, Michael Dyer developed a simulated environment containing stationary females that were capable of emitting a 3-bit "message" and mobile males that could "hear" the females' signals, but were blind. Reproduction could only occur when a male and female actually found each other. There was no partial reward for coming close. With only this binary intrinsic fitness criterion, the males and females eventually evolved a simple language in which a female would guide the nearest males to her location with a series of "instructions". The various evolutionary stages that this simulation went through on its way to the final language are an interesting lesson for those who find it difficult to see how complex problems could ever be solved by "just" natural selection.

Most evolutionary simulations, such as mine, utilize only extrinsic fitness criteria simple because of limited computational resources. Intrinsic fitness requires that a complete ecosystem be simulated for a long period of time with a large population (Dyer's work used a population size of 65,000 and was performed on a Connection Machine). Instead, I selected an external fitness critera which estimates the probability of survival given only a brief evaluation and a (relatively) small population size. The basic idea of choosing distance from the patch was that, on average, the agents that are found closer to the patch are more likely to eventually find the patch if given sufficient time than those that never come near it. Given that I was able to evolve true chemotaxis in this fashion suggests that the external criterion I chose was a good estimator of true fitness.

More generally, the point of my work with genetic algorithms is not to construct biologically accurate evolutionary simulations. Rather, I am interested in GAs as a (semi)automated design methodology for producing autonomous agents that are well-adapted to given environments. I am also interested in GAs because they allow me to generate controllers based only upon external requirements and therefore without any preconceptions about their internal organization. I think that examining the internal operation of such artificially evolved controllers may significantly expand our currently rather limited imaginations when it comes to the metaphors we use to understand biological controllers or to engineer artificial ones.

> [discussion of reorganization]

I agree that there are a great variety of "reorganizational" processes between motor control and evolution, though, given my comments above, I certainly wouldn't short-change natural selection nearly as much as you do. A number of simulations in recent years have demonstrated that an ability to learn can significantly speed evolution.

The specific reorganization process that you describe is an obvious choice for negative feedback systems, since an explicit error signal is conveniently available. Gradient descent, a common learning procedure in the neural network literature, would just be a special case of your biased random walk which guaranteed to reduce the error at each step. Your proposal would also seem to have strong ties to simulated annealing. As for negative feedback control, I think it likely that something like your proposed reorganization might occur in the nervous system, but I strongly doubt that it is a universal process for plasticity. There are many kinds of plasticity in the nervous system which do not have the benefit of an explicit error signal.

For example, the cockroach escape response that I briefly described in an earlier message is triggered by hundreds of wind-sensitive hairs that are found on the underside of two antennae-like structures known as cerci that are found at the rear of the animal. If one of these cerci is covered with wax or removed, the cockroach initially makes very bad turns, just as often turning toward the predator as away.

However, if it happens to survive for a period of roughly thirty days, the cockroach largely recovers the directionality of its turn.

Since the cockroach escape response is not organized as a negative feedback control system, it should come as no surprise that the mechanism underlying this plasticity is NOT any kind of error-correction process. Rather, the sensory inputs from the missing cercus atrophy, causing the inputs from the intact cercus to either grow more synapses or strengthen their effect in some other way (the precise biochemical process has not yet been worked out). The cockroach does not even have to actually escape many times for this reorganization to occur. The reorganization is an intrinsic cellular response to the loss of sensory inputs that has the appropriate behavioral effect. Since natural selection only selects for appropriate behavioral effect, that's all that matters.

Best Regards,
Randy Beer

Date: Mon Mar 16, 1992 2:02 pm PST
Subject: RE:blindfolding: Nevin

March 16, 1992

From: Bob Clark

Sorry (again) for the delay -- my life continues to become more complicated. Also, I have not yet caught on to all the operating aspects of working with the NET, and my CSG Files are not yet under control.

You "couldn't find the posts to which I was responding." There were five relating to Social Control. One from me, 2) your response Feb 10, 3) a comment from Powers, Feb 11 in which he mentions "blindfolding," 4) one from you "points of view" Feb 12, commenting on Powers "blindfolding," 5) the recent one from me of March 10. I expect that you have all of these except the one I added in my March 10 post: 'STATS VS CRUISE CONTROL, posted Jan 31. This is from a series I call "THRMSTTn."I don't have other identifying info -- I didn't realize more specific identifiers are available. I will post this Jan 31 item again shortly.

In response to your post of March 10 (is: "Tue 920110 14:18:16" the identifier?), you respond to my questions with your own: "How can one person model his behavior on that of others?" You continue: "If we understand that in a clear and precise way, I believe we thereby have clear and precise answers to your questions." I agree with you strongly. That is why I asked my questions. I think these matters need study and that answers can be developed. However I think it is necessary to review the basics of Hierarchical Control System Theory and make a few changes.

I hope my remarks about idealized "Black Boxes" (the ECS you explained to me -- thanks) can eventually lead to such refinements.

Since I have been out of touch for over 25 years, much has been done that I've not seen. I have Bill's important book, Robertson's book and the collection of Bill's papers. Reading between the lines of some of the papers on the CSGNET, I get the impression that certain key concepts need examination.

I expect to develop these ideas in the sequence I am calling "BLCKBOXn," posting the next entry soon.

Regards, Bob Clark

Date: Mon Mar 16, 1992 2:43 pm PST
Subject: 'STATS VS CRUISE CONT'L

FROM BOB CLARK

This post was a response to a suggestion from Gary Cziko.

Your suggestion about Cruise Control is interesting. It can lead to a discussion of several aspects of control system operation: sensitivity, response time, offset (or "dead zone"), relation to other systems, etc. These are interesting and important in themselves.

However I selected Thermostatic Systems because they are familiar to many people and they include the basic elements of Negative Feedback Control Systems. That is, they include: 1) a means for detecting a variable, 2) a means for affecting that same variable, 3) a means for subtracting the magnitude of that variable from some "preset" value, with a resulting positive difference acting to produce a positive output from the second "means". This is the usual combination of components composing a Negative Feedback Control System.

Such a system need not have a continuous output to achieve its result. It is interesting to observe that "continuity" is, in part, a matter of "viewpoint." Thus, if the Thermal System is observed over a period of several hours, its control approximates continuity. And the Cruise Control, observed in milliseconds, reveals various limitations.

Also, it is true that the Thermal System is a "one-way" system as usually presented with a furnace, etc, that is only one of several limitations it suffers. Another "one-way" system is a living muscle fiber! It can only pull, not push.

The Thermostatic System can also be used to illustrate other aspects of control systems -- and other forms of control system. Thus:

Before Thermostatic Systems were developed, people kept warm in the winter. My father had a coal fired furnace that had a damper that adjusted its operation. Too cold -- open the damper; too ward -- close the damper. And it was a fairly continuous operation.

Where was the Control System? Clearly the situation was livable, although not as convenient as one would like. Obviously, there was a control system in operation where the temperature (where?) was the controlled variable even though the control was accomplished by adjusting the rate of heating.

But where was the Control System? Without a person, the temperature was not controlled -- but also without a damper and a fire box,the temperature also was not controlled. Some person - DECIDES - whether action is needed, and in which direction. He then uses his (lower order) muscle systems to affect his - ENVIRONMENT - according to his - UNDERSTANDING - of his environment.

What about "his - ENVIRONMENT - ?" This usually refers, perhaps vaguely, to his physical surroundings outside his skin.

But someone else might be available and asked to "open the damper," "turn up the furnace," etc.

Where now is the Control System? The person could "do it himself" or "ask someone else." Having made his - DECISION - , he used his lower order systems to get his desired result. Were there two (or more?) levels of control involved: the "other person" and the "furnace."

Thermostatic Systems now take the place of the "other person," much more efficient and convenient.

For those who are familiar with thermostats, most of this is unnecessary. But what about those whose - ENVIRONMENT - does not include - UNDERSTANDING - of control systems? You must have seen people turn the setting up higher and higher when the room doesn't warm up fast enough? The furnace was already at full speed, so raising the setting has no immediate effect. Later, however, the room is too warm and the setting is reduced. This is "over-control" and the system is oscillating!

Notice the importance of the - TIME - scale of the person vs the "response time" of the System.

This illustrates the difference between regarding the assembly of parts as a Control System rather simply a group of connected parts. Such a difference in - VIEWPOINT - can result in a difference in behavior. Often these differences have little effect, but sometimes they are very important!

Both viewpoints are valid and result in the same mathematical representation. However one is more useful for using the system, the other for modifying the system.

Several words have been noted above: DECIDES, UNDERSTANDING, ENVIRONMENT, CHOICE, DECISION, VIEWPOINT, TIME. These words and their associated concepts are used routinely and seem to be readily accepted. However each of them is very important, and merits closer examination.

And how does each relate to a Hierarchy of Control Systems?

Regards, Bob Clark

Date: Mon Mar 16, 1992 11:05 pm PST

Subject: Alpha control; Genetic Algorithm; BEERBUGS

[From Bill Powers (920316.1600)]

Note to Bob Clark and others on identifiers. Various machines on this network supply gobs of useless header information (and mine puts more AFTER each message) in which it's hard to tell who the message is from. Some mail systems lop off all that info, or most of it, before you ever see it, and some don't. We have evolved a de facto semi standard by which we put INTO THE TEXT the information that's otherwise hard to find or missing. I always put

[From Bill Powers (yymmdd.hhmm)]

at the beginning of my text, as above, so the reader can see who this is from without searching to the end or figuring out header stuff. When the text is a reply to someone's post, I put a line like

Bob Clark (920316) --

just before the reply, sometimes with a ---------------- separator at the end.

Also, on quoted material that starts

>like this
>and this

... some systems allow you to import text from other files, and put those > signs in automatically. Mine doesn't. I copy the material from another window in my word processor, replace the hard returns with spaces, and insert the >s by hand.

All this stuff is typed in by hand -- it isn't a feature of the network.

Bruce Nevin (920315) --

Neural network simulators: thanks, but no. It's too hard to get the materials and learn them, and anyway I do enjoy reinventing wheels just in case something a little different turns up. From a couple of simulators I've seen, I couldn't run them on my machine and wouldn't much want to. If I find myself desperately in need of such things I will let you know with gratitude.

Rick Marken (920313) --

>So if the gamma efferent is fixed at some value won't it keep stretch
>signal (and, hence, load position) constant despite variations in alpha
>efferents changing the requested force exerted on the load?

The gimmick is that the alpha reference signal enters the motor neuron as a force reference signal. The stretch signal that enters the same neuron is a length ERROR signal (the comparator is mechanical) serving as another force REFERENCE signal. So if you put in an alpha signal, it's basically telling the stretch system that there's an error when there isn't one. As a result, the stretch system alters the position until the sum of its error signal and the alpha input signal is zero again. Of course that happens at a different position.

The main reason I use the alpha signal is that when the gamma reference signal changes in a large step, it creates a large error signal that is made still larger by the rate-of-change component that follows instantly. The model just doesn't seem able to handle the transients very well. All that may be just superstition, however. All I know is that when I switched to the alpha reference signal and did (probably) a lot of other things too, the model suddenly got much more stable. Go figure.

I've sent Arm Version 2 off to Greg Williams for evaluation and revision, and have started the writeup. I'm using Pat and Greg's PictureThis to make the diagrams, and liking it better all the time. I should now be able to get back to Bug programming, and when that's on a plateau I'll get back to the frequency tracker for sound spectrograms (you thought I'd forgot?) and when that's popped off the stack ...

Martin Taylor (920313.1100) --
Rick Marken (920313) --

>So if the gamma efferent is fixed at some value won't it keep stretch
>signal (and, hence, load position) constant despite variations in alpha
>efferents changing the requested force exerted on the load?

The gimmick is that the alpha reference signal enters the motor neuron as a force reference signal. The stretch signal that enters the same neuron is a length ERROR signal (the comparator is mechanical) serving as another force REFERENCE signal. So if you put in an alpha signal, it's basically telling the stretch system that there's an error when there isn't one. As a result, the stretch system alters the position until the sum of its error signal and the alpha input signal is zero again. Of course that happens at a different position.

The main reason I use the alpha signal is that when the gamma reference signal changes in a large step, it creates a large error signal that is made still larger by the rate-of-change component that follows instantly. The model just doesn't seem able to handle the transients very well. All that may be just superstition, however. All I know is that when I switched to the alpha reference signal and did (probably) a lot of other things too, the model suddenly got much more stable. Go figure.

I've sent Arm Version 2 off to Greg Williams for evaluation and revision, and have started the writeup. I'm using Pat and Greg's PictureThis to make the diagrams, and liking it better all the time. I should now be able to get back to Bug programming, and when that's on a plateau I'll get back to the frequency tracker for sound spectrograms (you thought I'd forgot?) and when that's popped off the stack ...

>Bill, I think you sell Genetic Algorithms (GAs) short.

I don't mean to. I'm just trying to indicate some areas where "reorganization" may be a better name, and a better model, for what may be overgeneralized as a "genetic" process. As its name and its ancestry imply, the genetic algorithm works by weeding out entire individuals simply by not allowing them to reproduce. I think this is a realistic model, and am prepared to admire the methods used in the modeling. I have no doubt that this method will create good evolutionary models. I seem to be getting a reputation as a GA-basher, which I would like to shed.

The place where I raised what I think is a legitimate red flag concerns the selection criteria, which are quite aside from mutation-vs-crossover questions. Randy Beer, in explaining the difference between internal and external survival criteria, has assured us that this problem is recognized, so I can stop worrying about it in that sense. I will still worry about it in the sense of wondering if the external selection criteria don't take us closer to a reorganizing-type process than a genetic-type process. But this is a tricky point and I don't expect it to be made easily. More on this in replies to Randy Beer, below.

Randy Beer (920316) --
Pardon the extraneous material, but most of this post is directed to you.

You're raising objections to the forcing of the control-system organization on everything just as a matter of principle. You're right, of course, and I can understand where that impression comes from. I do try to see the closed loop first, and give it up only after some trying. But I am willing to give it up when warranted (or at least back up a couple of steps). In part there's some confusion because I'm used to dealing with large organisms in which sensory feedback is of tremendous importance. Simple organisms don't necessarily need it for everything. But I am allowed to wonder whether a cockroach's turning "away" from a threat couldn't be interpreted, from the cockroach's frame of reference, as moving the perceived threat around to the rear and then making it as distant as possible. A sixty-millisecond action doesn't seem too short for feedback control. Even the human arm positioning system can move the forearm from one stable position to another in around 100 milliseconds under active control (although that's the lower limit and involves tremendous muscle forces -- 20g accelerations). In a cockroach where the path lengths are only a few percent of those in the human spinal systems, I should think the loop delays can't be more than 5 milliseconds, which allows plenty of leeway for performing a movement in 60 milliseconds under good feedback control.

Even if I have to back away from such detailed control systems on occasion, there is always the larger question: what effect does the action have on the animal? Between natural selection and larger control loops there is room for many answers. I guess I just want to bring this point of view more forcefully to attention, to compensate for the way biologists tend to dismiss it (usually for mythical reasons such as "feedback is too slow"). If more biologists actually understood the properties of control systems, I would take their objections more seriously.

>... it seems that you prefer to make everything purely sensory-driven.
>This is, of course, a perfectly legitimate approach, but I must repeat
>that it does not appear to be the way biology does it.

This illustrates part of the problem. Control systems are not "sensory driven." In fact the best interpretation is that the actions of the system drive the sensory information toward states specified inside the system: inputs, not outputs, are "controlled." The external effects we observe are side-effects of this process.

When you say that this does not appear to be the way biology does it, you're asserting a conclusion that I'm reluctant to take as a general principle. If you were to say that this isn't the way biologists interpret what they see, I would agree without hesitation. But since biologists have generally rejected control theory without understanding it, this leaves room for guessing that perhaps what they reject as an alternative explanation is not control theory, but a set of conclusions they incorrectly attribute to control theory.

For example, as mentioned, one basis for rejection is the idea that a feedback control system is slower than an open-loop system producing the same kind of output. This is the exact opposite of the truth in almost every case. A system with feedback can operate much faster to produce a given amplitude of movement for the simple reason that the unfedback gain can be far higher than in the open-loop system. If the open-loop system started with the same initial rise of output that the feedback system starts with, its final action would be grossly too much. To prevent overaction, the open-loop system must be run at appropriately low gain and correspondingly slow response. With feedback, the frequency response of the system becomes almost flat over a much wider range, meaning that slow actions and fast actions will reach the same amplitude. The first systematic use of negative feedback was to increase the bandwidth of vacuum-tube amplifiers (H.S. Black, in 1929).

Another basis for rejection is the myth that in nervous systems, time-delays make all strong negative feedback unstable. This too is untrue. In the first place, the delays usually mentioned are those associated with "reaction time," which occurs only for complex kinds of controlled variables which even real people control slowly. In the human spinal control loops the transport lag is around 9 milliseconds, not the usual quarter of a second that is mentioned. Also, few critics of control theory are familiar with simple stabilization methods that easily compensate for time-delays while allowing very high loop gains to be used. With proper filter design, the result can be full error correction within a single transit-time of a neural signal, or at most two.

Assuming that a cockroach's neural signals travel at only 1 meter per second (unmyelinated), a loop that's 5 millimeters around would have a transport lag of 5 milliseconds. This would put a frequency limit on the associated control system of around 100 Hz, if the leg dynamics allowed it. A 27 Hz running frequency would be quite feasible.

So when you say that biology doesn't do it that way, I hope you'll forgive me for wondering just what "way" you have in mind.

Your argument that there is a central pattern generator in walking, that can continue to work without sensory feedback, is convincing. In the design I'm working on I have incorporated a nonsensory feedback path that allows for this, and in fact is part of the pattern generator. I would be very surprised if the design I come up with eventually doesn't resemble many features of your model and those of others. Despite appearances, I'm trying to do the modeling honestly first, and according to PCT principles second.

On to second post:

>Many consequences that are of utmost importance to an animal are not
>controlled in any negative feedback way.

If the animal's actions in no way alter the effects of those consequences on the animal, then I would agree with you. But it's hard to imagine a behavior that has consequences that are "important" to the animal that the behavior doesn't do something to control. For instance, if there isn't enough of the consequence (food, mates, temperature), the behavior ought to do something to increase it; if there's too much (which could mean any at all), the behavior should do something to decrease it. That's control. The moth that folds its wings and drops like a stone is controlling the intensity of the bat's echolocation sound -- making it less, if the maneuver succeeds. That's important. But this little control loop makes possible a larger one: keeping the distance to bats as large as possible, a goal that other means also serve. And that makes possible a still larger control loop -- keeping from being eaten. The latter, of course is a loop that we call evolution (some don't think of as a loop at all).

You say of the cockroach's escape response, "people have found that it is organized not as a negative feedback system, but as a feedforward system." By this I presume you mean an open-loop system -- the response has no effect on the sensory inputs that lead to it. This, to me, is a very suspect "finding" -- it sounds more like an opinion. It's hard to imagine how a cockroach's body could begin to turn without instantly starting to alter the sensory signals that gave rise to the motion, or how moving away from an approaching object could avoid altering the sensed relative velocity of the object. How can you avoid sensory feedback in such a situation?

>The sole purpose of the escape response is to make sure that the
>cockroach isn't where it was when the attack began. Consequently, the
>cockroach can complete its initial turn in about 60 MSEC. Given the
>latencies involved in the sensory organs and neural signal >transmission,
there simply isn't time to do negative feedback control >of this turn.

I think this understates the specificity of the purpose, unless cockroaches run just as often toward the threatening object as away from it. I would say that the purpose is to reduce the proximity of the threatening object to as low a level as possible. This looks like a control system to me. The reference proximity is zero. The sensed proximity is nonzero. The action makes the sensed proximity approach the reference proximity. That's negative feedback no matter how you look at it. And see above remarks about those "latencies."

>The actual movements required to escape might be very different in each
>of these [mentioned] cases.

This is the hallmark of a control system: the action is whatever is required to make the sensed condition match the reference condition. The output is VARIED in order to control the INPUT. This result can be achieved by open-loop logic that takes every possibility into account, but it is achieved far more simply with control systems, which control outcomes by varying the means. You can build a thermostat that senses an open window, a fire in the fireplace, the outside air temperature, wind, and solar flux, and the number of people in the room, and compensates by adjusting the furnace output -- or you could just give it an air temperature sensor and ignore the reasons for the temperature changes. Maybe very simple organisms do some things open loop, because they are so simple and live in such simple environments that the same action always has the same consequence. But I am dubious about such claims.

>... do you seriously believe that an animal explicitly estimates its
>current survivability and that some high level control system actually
>uses the error between this estimate and a "reference" level to guide >its
behavior?

No. I don't think that organisms control for survival. They control for enough food (or enough nutrition), the right temperature, presence of mates at the right times, odor concentrations, and many other things. The consequence of controlling these variables is some degree of survivability -- in the eyes of the observer. Only an organism that can perceive in terms of survival can control for (perceived) survival (potential).

>What I was trying to say in the summary of my own position was that
>evolution selects for animals that always generate the appropriate
>consequences PERIOD.

I agree. It's not easy, however, to decide which of the many aspects of behavior one can observe is the appropriate consequence for the organism. It's also not easy to explain (without control theory) how a consequence can be selected for when a different behavior is required each time it's produced.

The cockroach's running has the consequence of moving its image on the retina of the observer. This, presumably, is not an "appropriate" consequence. But what of the consequence of moving its body by a certain path across the room toward a food-patch? What if it moves to the food-patch but fails to eat it? Eventually we get down to the consequences that matter: consequences that affect the internal state of the cockroach. The smell of the food affects the internal state. The avoidance of bright light affects the internal state. Eating the food affects the internal state. Fleeing from approaching boots affects the internal state. The only consequences that matter to evolution (or behavior) are those that affect the internal state of the cockroach. All other effects are side-effects, no matter how fascinating to the human observer.

In HPCT, we divide these effects on internal state into two classes: effects on the life-support machinery, and effects on sensory signals that respond to the environment. Evolutionary theory introduces a third effect: effects on reproduction. The effects of given external events on the capacity to reproduce, however, depend critically on the organism's ability to control the other two classes of effects.

In HPCT, reorganization is based on sensing the state of the internal biological machine. The sensing can be largely biochemical, or it can involve sensory systems in higher organisms (the autonomic system, for example). The reference levels for these internal variables are specified in DNA. The error signals drive random reorganization. The result is a behavioral system that can control the sensory effects of external events. When the sensory effects are adequately controlled, the internal state variables are maintained near the inherited reference levels, so reorganization goes at a slow or zero rate. When the reorganizing system is maintaining the internal state close to the inherited reference levels, the organism remains capable of surviving to reproduce, so the pressures that result in evolution are prevented from acting. The species evolves at a low or zero rate, because the optimum variants of intrinsic referenc levels are already present. Each tier of this three-tier process accounts for the next one up, and acts only until the next one up removes the pressure for change. There may be more than three tiers, but you get the idea.

Which brings us to your third post (I'm catching up after a weekend of a failed network node): the Genetic Algorithm.

>So there is a very small chance that even a very unfit individual will
>get to reproduce, though the vast majority of selected individuals will
>be very fit.

I trust that you mean "in the model" and not necessarily in real organisms. A population that is very unfit, to the extent that it can't manage to move any of its individuals to the food patch, will not actually survive -- the tail of the probability distribution can't make up for lack of a critical motor neuron. I think we have to be ultra-cautious about statistical arguments: they aren't constrained to make sense.

>Extrinsic fitness is when the experimenter imposes an external fitness
>criterion of his or her choosing. Intrinsic fitness is when the >fitness
of an agent is directly tied to its ability to find a mate and >reproduce
within the simulated environment.

Nice clarification. I understand why extrinsic fitness is sometimes used, just to see the effects of a model. But there's a deep difference between extrinsic fitness criteria and intrinsic ones. The extrinsic ones are actually changing the nature of the model, because they allow the detection of the direction of change without requiring an actual achievement of something that affects intrinsic fitness. An external direction of the course of evolution is introduced, through an intelligence that knows which changes (worthless in themselves) will lead to eventual structures that will enhance survival. I understand your argument that using only intrinsic criteria would require complex modeling and very tiny evolutionary steps. But that's the reality of the situation. I'm not convinced that what you get through extrinsic criteria is the same sort of thing you will get through intrinsic ones which have no sense of direction.

I'm trying to put forth another notion, called reorganization, that can accomplish the sorts of behavioral changes that are seen in your GA models. I think that when you fully grasp how this sort of system works, you will see more potential in it than you do now. My assumption that you don't grasp it yet is based on statements like this:

>Gradient descent, a common learning procedure in the neural network
>literature, would just be a special case of your biased random walk >which
guaranteed to reduce the error at each step.

The biased random walk involved in reorganization is NOT guaranteed to reduce the error at each step. In fact, at each step, the likelihood of decreasing the error is the same as the likelihood of increasing it. The random changes themselves are not biased in any direction. They are PURELY random. What introduces the bias is the frequency with which the random changes are made, which in turn depends on selection criteria. The statistics would be almost identical to evolutionary statistics in which mutations and crossovers are truly random and the result is biased by the selection effects.

It's hard to convince onself that such a process could ever have a systematic effect. Rick Marken and I developed a very simple demonstration that's worth setting up just to experience it. You put a dot on a screen that moves at constant velocity in some constant direction. You put a target position somewhere else on the screen. Then you arrange for each press of the space bar to alter the direction of motion AT RANDOM, using a random number generator for the angle. It's hard to believe, but just by pressing the space bar, you can steer that dot right to the target with far less wasted motion than you would intuit, if you believe that you could ever get there at all. Of course when you get TO the target the thing keeps moving so you begin to develop a messy track that weaves all around and through the target position -- this is the equilbrium condition of the control system.

The selection criterion is simply distance from the dot to the target, or better, a negative rate of change of distance. This looks like a two-dimensional situation, but in fact if you compute as if for two dimensions but display only the radial distance to the target, you will still get there just as easily. You will also get there (but you won't know it) if you display only the radial velocity toward the target (the projection of the two-dimensional velocity onto the radial direction).

To make this work for your chemotactic system, you need some selection criterion that will be a measure of approach toward the food-patch. You are, in fact, using one -- but it's not intrinsic. I would suggest trying odor intensity. There may be other criteria you can think up -- intrinsic criteria -- that will offer the right kinds of biases. Of course to play fair, these criteria have to be built in and unchangeable, not being altered by the organizations that result. With evolution as a source of criteria you can be generous. Maybe falling down could create some basic side-effect that's undesireable, and so on. It's a lot easier to think of evolution as providing these basic criteria than it is to think of it working out the details of neural circuitry appropriate to the current environment, including parameter values. With inner criteria to guide single-lifetime reorganization, each organism could try thousands of times to adjust each parameter, instead of different individuals getting only one tryout of a new organization per generation. The selection process from generation to generation would still be based on survival, but now what has to evolve is not the neural circuitry, but the criteria for reorganization. I have a hunch that this would prove to be a very powerful combination. It ought to handle the waxed-cercus problem easily. It would certainly handle the variability of neural circuitry from individual to individual -- with the same functions being accomplished -- better than simple from-scratch natural selection could.

God, eight pages. Do I dare send this? I guess the answer is yes, if you get it.

Best   Bill P.

Date: Tue Mar 17, 1992 11:13 am PST
Subject: mostly language

[From: Bruce Nevin (Tue 920317 09:14:28)]

New Interleaf software is being installed today, so my workstation is down and I have some time to catch up on some responses I've been wanting to make.

Randy Beer (3/16/92 12:04) talks about cockroach escape routines that appear to run open loop.

One perspective on emotion is that it invokes or is invoked together with quick, "canned" responses to stereotyped situations. The inaccuracy of the stereotyping (category level, I assume) is offset by the survival value of immediate fight/flight response.

The stereotyped situations can be social, involving but not necessarily limited to conspecifics. My daughters just went this weekend with friends to a place called Wolf Hollow, where they have a pack of wolves and teach the public about them. In teaching about social hierarchy, the woman demonstrated that she was lowest in the hierarchy: when she approached while they were feeding, every tail would rise, with growls. She explained that this is what happens if a wolf approaches before a higher-status wolf has finished eating, and the lower-status wolf's tail goes down in submission as she/he retreats. She said that her husband is #3 in the pack hierarchy, after the head male and head female, but could not demonstrate this because he was pumping a visiting animal ethologist. The threat to survival here is less immediate than an attacking predator, but real. Given more predictability of the other players, one can focus more attention (higher gain) on perceptions whose control contributes in more obvious ways to fitness and survival. I suspect this is why assessment and communication of relationship is universally of high importance among mammals, and probably lower orders as well.

Avery (3/4/92, 3/6/92) --
You speak of semantic representations for words, overlapping for pairs like betray and betrayal. You refer to Jackendoff's 1990 "conceptual structures." I have not read this, and am unlikely to given my situation. Please sketch what this looks like for the kinds of examples we have been discussing. Am I correct in assuming that it is the semantic "conceptual structures" that get correlated somehow with nonverbal perceptions in the perceptual hierarchy, rather than the words for which these semantic representations are provided?

I speak of semantic representations as well. For me, the semantic representations use only the words and word relationships that are characteristic of ordinary language. For the sake of refining the semantic representations of discourses, the method extends some word relationships beyond the domains that are normal (conventional, natural) for ordinary, everyday usage of language.

>In effect, rather than
>`expand' whole sentences as syntactic structures, one `expands'
>lexical items into their semantic structures, plugging arguments
>into appropriate positions

Since the reductions take place at the time of word entry this is essentially what I am doing. See more detail in my response to Bill later in this post.

(3/5/92) --

>On `Jimmy's betrayal': why do you think I or any other Chomskyan
>needs a [+abstract] feature here? All I think I need is (a)

A recent query on the Linguist Digest used semantic features of this sort, viz. Michael Newman <MNEHC@CUNYVM.bitnet>, posted 3/6/92 and appearing in Linguist Digest 3.233 of 3/9/92:

>The problem is related to the semantic distinction between referents
>which are +specific (or +referential) or more informally real, existing
>out there in some way, and those which are -specific, hypothetical or
>generic. Thus there is a clear distinction between two readings of . . .
>"Peter is going to marry the richest woman in town." So far so good. The
>issue seems relatively clear when you are using example sentences, but
>when you use real language, things do get messy sometimes. For
>example, I would be reluctant to use the + or - specific label for the
>following example from my corpus, (which is based on TV talk shows by
>the way) The problematic antecedent-anaphor pair are in caps:
>
>   (1) I have become involved with a consumer advocacy group called s.h.a.m.e.
>   it stands for Stop[Hospital and Medical Errors, and it is a group that
>   was formed by MALPRACTICE VICTIMS and THEIR families.
>
>In this case there were indeed a concrete set of people who formed this
>group, yet neither speaker nor hearer were in any position to specify
>that set any further. In addition it is certainly conceivible that there
>might be some dispute as to who exactly belongs in this set or not. So
>my solution was to label this type as semi-specific (actually I use the
>term 'semi-solid' reserving 'solid' for specific and non-solid for
>-specific, but I don't want to get into that here) Cases like (1) where
>there are sets which none of the interlocutors are in any position to
>identify are fairly common, and using this semi label, I have managed to
>reduce the number of problematic tokens by more than half.

Words like abstract, specific, semi-specific, solid, etc. are words of English and depend upon the background vernacular of ordinary English for us to interpret and use them. Yet the claim in such systems of semantic representation so far as I am aware is that they are a universal vocabulary of a universal semantic metalanguage apart from English (or some other language being described). This metalanguage has its own syntax involving constructions such as trees and tables. Generally not noticed is the fact that this semantic metalanguage in turn requires a semantic interpretation, and that its seeming explanatory power depends in a circular manner upon the background vernacular of natural language. Recourse to nonverbal perceptions in the perceptual hierarchy as the universe of "meanings" promises a way out of this circularity. The question remains: if the work can be done within language (with slight extensions to enhance regularity), why have recourse to a metalanguage purportedly outside the familiar language of words organized in sentences and texts?

(Bill Powers (920310.1700)) --

>Here's my diagram of the operational definition of a controlled variable,
>with those "low correlations" shown:

```
>
>                              ------- low correlation ------
>                             |                               |
>       ------------          |      -------------        --------
>      |            |         |     |             |      |        |
```

```
>      |  BEHAVING  |----->ACTION-->|ENVIRONMENTAL|----->|OBSERVED|
>      |   SYSTEM   |               | TRANSFORMA- |      |OUTCOME |
>      |            |               |    TIONS    |      |        |
>       ------------                 --------------       --------
>                                                          |  ^
>                                       --low correlation--   |
>                                      |                      |
>                                      |                      |
>                         VARYING INDEPENDENT DISTURBANCE -->--
```

Is there not a high negative correlation between the action and the disturbance, both measured in units of effect on the observed outcome? And is it not that we infer from this, on the basis of coherence of theoretical explanation and on the basis of our introspections, that the behaving system has within itself (hidden from us) a perception of a goal, the reference perception, and that the above high negative correlation is a pretty exact mirror of a high positive correlation between the goal and the outcome?

I suggested (3/9 midday--I didn't finish putting the date in!) that correlation of an outcome with a purpose or goal smells unscientific from a conventional perspective because the goal is a perception hidden within the black box of the behaving system and the controlled outcome is that subset of the possible sensory inputs to the behaving system that are relevant to that goal.

I understand that we tease out which aspects of the outcome are controlled and which are accidental byproducts by observing whether or not the behaving system resists disturbance to particular perceptual inputs (taken from the perspective of the behaving system). To undertake this surprisingly challenging work, poking around in the dark far from the seeming light under the lamppost, one must first have grasped and at least suspended disbelief in the theory that both motivates and guides it. Thus:

>This invariant, the observable outcome, is not directly observable as a
>first-order observation. It is a second-order observation, observable
>only in context of certain expectations derived from hierarchical
>perceptual control theory.

I was concerned with what might constitute stumbling blocks in getting an adherent of the traditional perspective to leave off stumbling around under that lamppost with its burnt-out bulb and come look where the key is more likely to be found. I tried to identify "unfamiliar steps" including:

> isolating an "invariant" outcome in the
>organism/environement system, selecting those aspects of the outcome
>that are relevant from the point of view of the organism, and applying
>the Test to verify that your guess as to what is being controlled is
>correct.

Have I missed an essential point, as you suggest (920310.1700):

>You guys are going all around the point here. I have a distinct feeling
>that you're avoiding it.

In your reiteration of basics, you say:

>Ordinarily, we would expect action and disturbance to correlate
>with the outcome.

Do you mean the algebraic sum of action and disturbance correlates with the outcome, while neither correlates well, taken separately?

>When there is control, the variance of the observed outcome is
>significantly less than the sum of the variances of the action and the
>disturbance.

I assume we are looking at the summation of variances of action and disturbance. Don't they cancel when there is control, but not otherwise? Doesn't the algebraic sum of variances of action and variances of disturbance approach identity with the variance of the controlled outcome, more or less nearly depending upon gain? I don't understand how the variance of the observed outcome can be significantly less than this sum of the variances of disturbance and action. What am I misunderstanding here?

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

I should say that I have never been invested in S-R models and have always deeply distrusted them, and that my field, linguistics, as I have learned it has not depended on S-R models. Even famous forbear Leonard Bloomfield, who is often caricatured in conventional histories of the field as a behaviorist, did not embrace S-R psychology so much as say (in 1933, abandoning Wundt for Watson) "this is as much as scientific psychology can tell us about meaning, which is not much, so let's put it over here in a black box labelled "semantics" and get on with what language can tell us in its formal structure, which is a lot." What I have of an S-R perspective must be in unconscious preconceptions taken in as part of the general American cultural package, but linguistics has been pretty free of it. Really. Chomsky's famous, withering review of Skinner's _Verbal Behavior_ was helpful to linguists who as in Bloomfield's day felt some obligation to go along with what the psychologists had to tell them, but met with no resistance to speak of within the field of linguistics, only among (some) psychologists.

Bill Powers (920311.0930) --

>>If you accomplish the aim of accounting for what all languages have in
>>common, and you show that it all comes down to characteristics of the
>>world of nonverbal perception plus fundamentals of physics and
>>chemistry in the environment, like the acoustics of the vocal tract--
>>having reached the state where linguistic universals are trivially
>>deduced from first principles, what would remain?
>
>Nothing. I think you're pulling back from reductionism, which isn't implied

No, I was trying to get at something else. There are some aspects of language that appear to be universal. There are other aspects that are defined by social convention. These conventional aspects of language are not universal.

The means in the perceptual control hierarchy for learning, maintaining, and orienting one's behavioral outputs to social conventions or norms are presumably universal. Those universal means may impose some universal characteristics or universal constraints on what is a possible social convention of language. Some language norms may be widely shared among different languages through the historical contingencies of language change (related languages) and contact ("genetically" unrelated languages, to use the standard metaphor). But the conventional aspects of language are not universal.

And having reached the desired state "state where linguistic universals are trivially deduced from first principles, PCT still has to describe or account for those remaining aspects of language that are socially inherited and which do NOT follow from first principles. From the point of view of any universal theory, they are arbitrary. We say horse instead of Pferd or jahhom, we say went instead of goed, we say I ate the fish instead of I the fish ate or ate I the fish, and all the rest, because of historical contingencies that must ultimately be taken as irreducibly arbitrary facts of social convention.

>How can you tell when you have a satisfactory expansion?

It uses at each point of word entry (operator entry) only the relatively small stock of words and word relations that are attested in many perfectly ordinary sentences of the language. It uses at each such point only reductions of word shape that are attested as minimal sentence-differences between pairs of perfectly ordinary sentences, such that the sentence-pairs informally meet my judgment of saying the same thing and as a check of validity meet the stated formal criterion for transformation (e.g. preservation of acceptability-difference over a graded set of such sentence-pairs).

>>In particular, I believe that operator grammar shows a simple structure
>>for language--a structure of word dependencies--that is universal and
>>that accords well with perceptual control,...
>
>I agree that it does, although you will have to agree that it doesn't
>completely fit natural language as it is spoken without introducing some
>important invisible processes which are in principle unverifiable.

The departures from ordinary usage involve minimally extending the domain of well-established reductions or other word relations, for the sake of attaining a more regular semantic representation. By "minimal" I mean the least amount necessary to attain that aim. By "regular" I mean such that each difference of form correlates with one and only one characteristic difference of meaning.

The discussion of "expansions" suggests that the semantic representation is a string of words that results from all the reductions being "expanded". "That which is a product of one discussing something which is a result of one expanding something . . ." for the first four words of the preceding sentence, for example. This is not the case. In the typical case (see note * below on "typical), a sentence occurs with other sentences in a discourse (text). It is the nonlinear structure of the discourse that constitutes the semantic representation. To bring out that structure in its most regular form (see above), one changes the form of most or all of the sentences so that they are all instances of sequences of word classes that recur through the discourse. This may involve undoing some but not all reductions, or undoing some and replacing them with others. (For the unreduced semantic primitives (words) are still present in the reduced form of a sentence, albeit in the form of affixes or even in zero form.) Strings that appear to be constructions of many words may turn out to be single "words" in the sublanguage grammar used for discourses in a particular subject matter.

Operator grammar is not the end product. It is a tool for analysis of discourse. A semantic representation for a discourse is an end product of that analysis. The word-classes of a sublanguage grammar correlate with high-level category perceptions for that subject-matter domain (e.g. symptom, patient, drug, etc. in a sublanguage of pharamacology.) The "words" that are members of these word-classes may look like phrases comprising a number of words for the grammar for some other domain. ("The beating of the heart" is the example I have given previously, a member of the "symptom" class for pharmacology.)

Does this help to alleviate some of your discomfort with the reductions of operator grammar? They are tools for changing the form of sentences to a more regular form relative to other sentences of a sublanguage exemplified by a set of discourses in the same subject matter, and for achieving a semantic representation of those discourses that correlates 1-1 with nonverbal perceptions of one conversant with that subject-matter domain.

A person's knowledge of a domain can be expressed in a set of discourses about the domain, or in a set of semantic representations (double arrays) for those discourses, or in the union of those semantic representations, or in the memory and imagination of nonverbal perceptions corresponding to the words and word classes in these. (* Note on "typical," above: It is in this context of knowledge that the atypical case of a single sentence or sentence fragment is interpreted. Hence the discussion of zeroed dictionary sentences, etc.)

>So I ask both you and Avery: in your models of language structure, which
>parts of the phenomenon of language are observed, and which are imagined in
>order to make the analysis work?

The words and word relations (including reductions) are observed, not imagined. Some extensions of word relations beyond their observed domains could be said to be imagined, though the bases for the extensions are well within attested variation for language. The regularity achieved by these extensions, and which motivates them, is observed, not imagined. The correlation of elements, relations, etc. in regularized discourses with nonverbal perceptions is observed. Even departures from regularity, and the development of different regularities replacing previously established regularities (as knowledge in a field changes over time) have their directly observable interpretation in the world of nonverbal perception.

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

(Bill Powers 920311.1100) to Martin --

>Suppose you want to produce a sequence like "now is the time for all good
>men to come to the aid of their country." By the wasteful pandemonium

>postulate, this is the province of just one sequence-recognizer and control
>system.

There might be one sequence recognizer for a familiar quotation like this one, as for a proverb, an idiom, or some other fixed expression. But in general, there cannot be a single sequence recognizer for each possible sentence in a language.

I think there must be a sequence detector for each operator word, satisfied by any word or words that meet the argument requirement of the operator, modulo reductions of those words, and within the bounds of a sentence (or sentence fragment) defined by intonation or punctuation. This intonation or punctuation is taken as a reduction of "I say" as highest (last-entering) operator (or "I ask," etc.).

That's all the news that fits, for now. Don't know when I'll have an opportunity like this again.

Be well,   Bruce

Date: Tue Mar 17, 1992 11:34 am PST
Subject: Reply to Beer

[From Rick Marken (920317)]

Well, I guess things have been as quiet as they seemed. I'll let Bill respond to Martin's interesting claims about GAs -- if he wants. I'll try to respond to the Beer posts. I'm glad to see that Randy is hangin' in there. I'm sure that he will find his attempts to get us to understand his point of view just as frustrating as we find our attempts to get him (and just about everyone else) to understand ours. I doubt that we will convert Randy -- but its nice to have someone participating on CSGNet who has such a pure understanding of the conventional point of view of behavior.

To begin:

Randy (to Powers):

>I continue to encourage your interest in models of insect locomotion.
>However, it seems that you prefer to make everything purely
>sensory-driven.

No, we prefer to make purposive behaviors have outputs that drive sensory inputs to internally specified reference levels. We don't think that those sensory inputs are really smart enough to drive anything anywhere but in random directions.

>This is, of course, a perfectly legitimate approach,
>but I must repeat that it does not appear to be the way biology does
>it.

Well, if "it" is purposive behavior then I think it is biologists, not biology, that is mistaken.

>A number of experiments have demonstrated that the neural circuits

>underlying many rhythmic behaviors (e.g. walking, swimming, chewing,
>breathing) can generate the basic oscillatory pattern IN THE COMPLETE
>ABSENCE OF SENSORY FEEDBACK.

Sure it can. But can it CONTROL the pattern (if that is what is controlled). If the "basic oscillatory pattern" is not a controlled variable then there is no reason to sense it. Did anyone test to see if the pattern is controlled? See my comments on "the test for the controlled variable" below.

> Of course, this central rhythm must be
>reinforced and fine-tuned by sensory feedback in order to exhibit
>completely normal output patterns.

I think we're talking "lower level control systems" here. And the "normal output patterns" suggest that these patterns are under control. The words "reinforcing" and "fine tuning" imply roles for these sensory inputs that they could not possibly play and result in control. If sensory inputs are in the control loop then they are controlled -- they don't "strengthen" or "guide" output.

> In addition, work on cockroach
>locomotion by Sasha Zill has suggested that even when sensory feedback
>is intact, it may come in too slowly to play any role in fast walking
>insects (the cockroach is capable of stepping frequencies in excess of
>24 Hz!).

This "speed of feedback" stuff is one of the BASIC MISCONCEPTIONS that led psychology (and biology and all the other life sciences) away from an understanding of the nature of feedback control in behavior. It is based on an S-R conception of how a feedback control system works. The idea is that a stimulus causes a response that has some sensory consequence (feedback) -- but if that sensory consequence doesn't register fast enough then it can't be of any use. But feedback loops don't work that way -- there is a continuous loop and feedback is ALWAYS their -- feedback is what is controlled. There are dynamic contraints on the operations of the loop --slowing factors and transport lags. But there are ways of dealing with transport lags (which is what Sacha thinks are too long) when they are long relative to the bandwidth disturbances to the controlled variable and it is also highly unlikely that the transport lag in the leg position control system of a roach is anything close to 60msec, it's probably about 1/10 of that -- plenty fast for control (if position is controlled).

In response to my query:

>>Would the acceleration of the bug as it falls off a ledge count as a>
>>behavior to be modeled? If not, why not. If so, why so?

Randy says:

>It certainly could. Some species of moths are preyed upon by bats
>who, as we all know, navigate by echolocation. These moths have
>evolved an interesting escape mechanism. Whenever they detect
>vibrations of a certain frequency (namely, that used by bats searching
>for prey), they simply fold up their wings and drop like a stone.
>.This certainly counts as a behavior in the ethological sense, though I

>couldn't say whether it is a behavior in the specialized technical
>sense of PCT.

You would have to test to see if the acceleration is controlled. I believe that the acceleration is simply an output that is part of the loop that controls another variable -- sensed intensity of sound at a particular frequency. This answers your question about what might be the controlled variable. The point I was making is that the conventional approach to behavior makes no distinction between controlled and uncontrolled consequences of neural output. This confuses the modeling process; it makes it impossible to tell when you need a control model vs a response generation model. I'm afraid that this confusion, combined with a strong bias toward output-generation models,has made it impossible for conventional life scientists to understand that purposeful behavior is the control of perceptual variables. The problem is that control can look like generated output on casual inspection; the study of "behaivor" must start with the test for controlled variables. Otherwise, you don't really even know what you are modeling. The "test" is described in Powers' Behavior: The control of perception. It is very important to understand the test. Without it, you have no way of knowing what an organism is doing (see my article "Behavior is the first degree" in the book "Volitional Action" edited by Wayne Hershberger).

>The basic idea of CT seems to be that behavior is the consequence of
>negative feedback control of selected sensory inputs.

Well, er, sort of. Controlled variables are behaviors that are a consequence of negative feedback control. It is also important to note that many variables that we call behavior (such as the spatial position of e. coli) are NOT controlled -- they are side effects of the control of other variables (for e.coli spatial position is not controlled -- just the perceived gradient of certain chemicals -- a unidimensional variable, not a three dimensional spatial variable). A control theorist would not call these behaviors "behavior". We would call them "irrelevant side effects"(IREs). I think many models of behavior are, in fact, models of IREs.

>However, if CT is making the much stronger claim that negative
>feedback control is universal and ALL behavior can be understood in
>its terms, then I am extremely skeptical.

No - the problem is the word "behavior" again. PCT says that ALL PURPOSEFUL BEHAVIOR (controlled variables) is the result of closed loop, negative feedback control of perceptual inputs. Other variables that might qualify as behavior but that are not controlled (not purposeful) are IREs --generated by the good old cause-effect processes that watchmakers have been familiar with for centuries. Cause-effect models don't work with controlled variables; you need to understand circular causality to understand control.

> Many consequences that
>are of utmost importance to an animal are not controlled in any
>negative feedback way.

Well, there are ways of testing this -- the test for the controlled variable. I think the "cockroach escape response" example of "open loop" control is actually a good example of a system busy controlling something very important to it -- its perception of forces, smells, etc that might be from predators (though all the cockroach knows is the percpetual veriables -- not "predator" -- that's our perception).

Finally, PCT has no organizational preconceptions. We are not interested in modeling a phenomenon with a control model until we know that what the system is doing is controlling. I would suggest that your approach actually does begin with an observational preconception -- it is assumed that "behavior" is any visible (to you, the observer) consequence of an organisms actions (where "actions" are any output of the organism that can be seens as a cause of the consequence called "behavior" -- ie. it could be a neural impulse, muscle tension or limb movement, etc). PCT starts by
2-40
recognizing that their are two different kinds of visible consequences of actions -- those that are controlled (produced on purpose) and those that are NOT (the IREs). We are aware of both kinds but only apply control models (closed loop, negative feedback systems) to the former.

Regards   Rick

Date: Tue Mar 17, 1992 8:07 pm PST
Subject: semantics, language

[From: Avery Andrews (920318:130615)

Bruce Nevin (920317.0914.28)

A Jackendovian semantic structure is basically just a sentence with explicit & unambiguous head-argument relationships, plus some semantic typing. `John betrayed Mary' might look like:

```
    [event BETRAY([person JOHN], [person MARY])]
```

where the BETRAY should be seen as an abbreviation for or pointer to a semantic decomposition of the meaning of the word into more primitive elements, and the lowercase words after the brackets indicate the ontological categories of the constitutents.

The verb betray might have a lexical entry like this:

```
  Category: V
  Form:     /betre:/
  Meaning: [event Does-sth-to(X,Y)
                  Y falls into the power of Y's enemies]
```

(J has a fairly elaborate theory of verbal lexical entries, which I'm fudging here_.

the noun `betrayal' would have the same meaning, but a different Form & Category, leading to differences in how the arguments get expressed.

I agree with a lot of what you say about semantic metalanguages, but it seems to me that there's a lot to be said for at least including fully explicit indication of coreference, predicate-argument structure dependencies, etc. One of the tasks I see of semantics is in fact to explain inferenceing abilities (the ones we actually have, like simple cases of modus ponens and universal instantiation), and I don't see how this is possible without adding the above adornments to the semantic metalanguage (I think that some of the Wierzbickians around here are beginning to get some of this message, at last).

Jackedendoff also wants to get an explicit representation of various patterns he sees, such as oppositions in the system of spatial concepts expressed by prepositions:

```
    in     into     out of
    on     onto     off of
```

He would certainly agree that the components of the semantic representational system have to be grounded in perception - actually doing this is one of his major interests, especially for prepositions, and has published some stuff with Barbara Landau on this (I can't recall where off hand).

I see two major weaknesses in his approach. The first is that he has nothing much to say about how linguistically derived information actually gets integrated from what is coming from the senses, and what is available from memory. I don't see any problems of principle here, and assume that he just hasn't gotten that far yet.

The second is that he seems to underplay the social aspect of meaning. For example, we on CSGNet can all refer to & say true things about plutonium, but I doubt that many of us have perceptual abilities that would enable us to distinguish it from enriched uranium without killing ourselves (well, I'm sure Bill could figure out how to do it, but I bet he'd spend a lot of time in the library before trying!!). The standard philosophical story about how this is possible is that we can talk about plutonium because we are properly plugged into a social system that contains people who actually can tell plutonium from other things, and this story sounds right to me.

This second weakness implies that at least some words will contain `irreducibly symbolic' components that actually can't be cashed out in terms of perceptions, or at least perceptions of normal speakers (non-experts). I suspect that even in the minds of real experts about solid subjects, theoretical terms aren't really perceptually grounded either. What is the perceptual grounding of `electron'? What is perceptually grounded is the description of various experimental setups and their outcomes, the best explanation for which is the existence of electrons. (This puts me in bed with those philosophers of science who think there really is a difference between language-of-theory and language-of-observational after all). So I see the term `electron' as having a rather complicated and indirect connection to perception, & suspect that it would be wrong to try to `expand' it into anything like perceptual terms. I'd sort of like to try to go on about this at greater length sometimes, but no more time, now.

If I get around to reading a reasonable amount of recent Harris, perhaps I'll insist that you read some Jackendoff!!

Avery.Andrews@anu.edu.au

Date: Tue Mar 17, 1992 9:17 pm PST
Subject: plummeting moths

How does describing the moth's `fold up & drop like a stone' routine as a feedback system amount to anything more than a decision to regard S-R systems as degenerate cases of ECS's?

Avery.Andrews@anu.edu.au

Date: Wed Mar 18, 1992 6:50 am PST
Subject: Principia Cybernetica Symposium-CFP

-------------------------- Original Message --------------------------
Call For Papers

```
*********************************************************
*   SYMPOSIUM:  THE PRINCIPIA CYBERNETICA PROJECT       *
*      computer-supported cooperative development       *
*        of an evolutionary-systemic philosophy         *
*********************************************************
```

as part of the

13th International Congress on Cybernetics
NAMUR (Belgium), August 24-28, 1992

About the Principia Cybernetica Project
____

The Principia Cybernetica Project (PCP) is a collaborative attempt to develop a complete and consistent cybernetic philosophy. Such a philosophical system should arise from a transdisciplinary unification and foundation of the domain of Systems Theory and Cybernetics. Similar to the metamathematical character of Whitehead and Russell's "Principia Mathematica", PCP is meta-cybernetical in that we intend to use cybernetic tools and methods to analyze and develop cybernetic theory.

These include the computer-based tools of hypertext, electronic mail, and knowledge structuring software. They are meant to support the process of collaborative theory-building by a variety of contributors, with different backgrounds and living in different parts of the world.

As its name implies, PCP will focus on the clarification of fundamental concepts and principles of the cybernetics and systems domain. Concepts include: Complexity, Information, System, Freedom, Control, Self-organization, Emergence, etc. Principles include the Laws of Requisite Variety, of Requisite Hierarchy, and of Regulatory Models.

The PCP philosophical system is seen as a clearly thought out and well-formulated, global "world view", integrating the different domains of knowledge and experience. It should provide an answer to the basic questions: "Who am I? Where do I come from? Where am I going to?". The PCP philosophy is systemic and evolutionary, based on the spontaneous emergence of higher levels of organization or control (metasystem transitions) through blind variation and natural selection. It includes:

a) a metaphysics, based on processes or actions as ontological primitives,

b) an epistemology, which understands knowledge as constructed by the subject, but undergoing selection by the environment;

c) an ethics, with survival and the continuance of the process of evolution as supreme values.

PCP is to be developed as a dynamic, multi-dimensional conceptual network. The basic architecture consists of nodes, containing expositions and definitions of concepts, connected by links, representing the associations that exist between the concepts. Both nodes and links can belong to different types, expressing different semantic and practical categories.

Philosophy and implementation of PCP are united by their common framework based on cybernetical and evolutionary principles: the computer-support system is intended to amplify the spontaneous development of knowledge which forms the main theme of the philosophy.

PCP is managed by a board of editors (presently V. Turchin [CUNY, New York], C. Joslyn [NASA and SUNY Binghamton] and F. Heylighen [Free Univ. of Brussels]). Contributors are kept informed through the Principia Cybernetica Newsletter, distributed in print and by email, and the PRNCYB-L electronic discussion group, administered by C. Joslyn (for subscription, contact him at cjoslyn@bingvaxu.cc.binghamton.edu). Further activities of PCP are publications in journals or books, and the organization of meetings or symposia. For more information, contact F. Heylighen at the address below.

About the Symposium

After the succesful organization of a symposium on "Cybernetics and Human Values" at the 8th World Congress of Systems and Cybernetics (New York, June 1990), and of the "1st Workshop of the Principia Cybernetica Project" (Brussels, July 1991), the third official activity of the Principia Cybernetica Project will be a Symposium held at the 13th Int. Congress on Cybernetics.

The informal symposium will allow researchers potentially interested in contributing the Project to meet. The emphasis will be on discussion, rather than on formal presentation. Contributors are encouraged to read some of the available texts on the PCP in order to get acquainted with the main issues (Newsletter available on request from the Symposium Chairman).

Papers can be submitted on one or several of the following topics:

The Principia Cybernetica Project
Cybernetic Concepts and Principles
Evolutionary Philosophy

Knowledge Development
Computer-Support Systems for Collaborative Theory Building

About the Congress
____
The International Congresses on Cybernetics are organized triannually (since 1956) by the Intern. Association of Cybernetics (IAC), whose founding members include W.R. Ashby, S. Beer and G. Pask. The 13th Congress takes place in the "Institut d'Informatique, Facultes Universitaires Notre-Dame de la Paix, 21 rue Grandgagnage, B-5000 Namur, Belgium". The official congress languages are English and French.

Namur is a quiet little city on the confluence of the Meuse and Sambre rivers, at the foot of a hill supporting impressive medieval fortifications. The congress atmosphere is relaxed and informal, with a lot of small symposia going on in parallel in adjacent rooms. There will be a welcome cocktail, a congress dinner, and a meeting room available for coffee breaks. Participants will receive a list of nearby hotels after sending in the registration form. They can also reserve inexpensive accommodation in student rooms.

Registration fee :

members of the IAC and authors of papers:	6000 BF (about $180)
other participants:	10000 BF (about $300)
Young researchers under 30 years (with certificate of their university)	2000 BF (about $60)

The fee covers congress attendance, conference abstracts and coffee-breaks.

Partial Congress Programme
____
The Congress will feature over 30 symposia, including the following:
(CHAIRPERSON Subject)

ACALUGARITEI G. (Roumania)
Evolutions and Metaevolutions from the Point of View of the Invariants Associated to the Transformation Groups

BAHG C. (China)
Complex Systems and their Evolution

COLLOT F-C. (France)
Les notions de temps et d' e'volution en Cyberne'tique

FRANCOIS C. (Argentina)
Les syst`emes humains home'ostatiques ou e'mergents

HEYLIGHEN F. (Belgium)
The Principia Cybernetica Project : Computer-supported Cooperative Development of an Evolutionary-systemic Philosophy

JDANKO A. (Israel)
- Cybernetic Systems Approach to History
- Cybernetic Systems Interpretation of the Religious Idea : From the Primitive to the Monotheist

GASPARSKI W. (Poland)
Cybernetics and Human Behaviour

GELEPITHIS P.(United Kingdom)
Invariants of Cognitive Science : Scope, Limits, Implications

STEG D. (USA)
Determinacy and Indeterminacy in Complex Systems

VANDAMME F. (Belgium)
Cognitive Modelling for Knowledge and Information Technology : Manual
and Automatic Tools

Balance deleted - Dag

Date: Wed Mar 18, 1992 9:17 am PST
Subject: plummeting moths

[From Rick Marken (920318)]

Avery Andrews (920317) asks:

>How does describing the moth's `fold up & drop like a stone' routine as
>a feedback system amount to anything more than a decision to regard S-R
>systems as degenerate cases of ECS's?

It's not just a matter of describing it as a feedback system -- it IS a feedback system. The best answer to your question is Bill Powers' article "Quantitative analysis of purposive systems" article in Psych Review (1978). The main point relevant to your question is that the moth's "sound control system" is just one of many cases where behavioral scientists have made the mistake of applying an SR analysis to a feedback control system. This mistake is even made in my field where psychologists doing tracking tasks -- and who know control theory -- describe these tasks in SR terms (for example, the distance from cursor to target is seen as the stimulus that causes "control action" responses). In the case of the moth, the sound emitted by the bat is probably seen as a stimulus for the drop response. Of course, since dropping will occur when anything produces the proper sound, the ethologists probably think of the sensed sound as the stimulus for the drop. Since it is likely that sensed sound is the controlled variable (with a fixed reference of 0) then the ethologists analysis of the moth is like the psychologists analysis of tracking (although, with the moth, the target is not "outside" as part of the stimulus as it is in tracking -- but since the moth's reference level for sound is probably 0, a sound level meter, which measures sound pressure relative to a reference 0 sound level- gives a measure of the "stimulus" which is like discrepency from the target in tracking).

So what is the problem with an SR analysis of the moth's control of sound intensity? First, I should say that, with a reference fixed at 0 a control model "looks like" an SR model -- even though it is not an SR model. There really is no way to model the moth's behavior as an SR model -- because there is also an R-S connection -- the sensed level of sound influences the dropping (S-R) but the dropping also influences the sensed level of sound (R-S). So there is a loop and any model of this process must take into account the dynamic constraints that make the control loop work -- that's one of the things in Bill's article cited above; he shows that a sequential approach (implied by S-R type models) to the moth's behavior (sound causes drop which then causes lower sound which then causes no drop, etc) will only work when their is very low gain in the loop; other-wise, the system become unstable.

So there really is no SR model of control (and the moth is controlling something). I have read articles about SR models of movement in simulated bugs -- the bug senses "light" and has outputs that move it toward the light. The people who build these bugs think of the bug as an SR device -- sensed light (S) causes response output (R). These bugs work, however, because the appropriate dynamics have been built in "accidentally" -- the programmers were not trying to stabilize the control system. For example, the effect of light on output is to produce increments in output proportional to the input. This incremental approach makes the "SR" device into a proportional control system with integrated output and a fixed reference level (which is implicit in the equations that transform input to output).

Most SR analyses of behavior (like those of the moth) are usually done in one's head -- it looks like SR. When working models are actually built (and they work) they are actually control systems with fixed reference levels (usually implicit in the S-R equations and, therefore, at the 0 point of the range of the input variable). These models are called SR models (the Braitenberg (sp?) "Vehicles" are another example that I just thought of) but they are NOT.

So, what's the problem with looking at behavior like the moth's in SR terms? Well, besides the fact that it's wrong and won't lead to the correct detailed model of the moth's behavior in a real environment, there are these explicit problems:

1) you won't realize that the reference for the controlled variable (probably the intensity of sound in a particular frequency region) can be changed (maybe) by the moth in order to accomplish some other goal (like, get around a tree -- I dunno). Anyway, SR analysis of control has this HUGE flaw -- it doesn't notice that variables are controlled; and possibly at varying reference levels.

2) you don't see the bat as just one of many disturbances to the controlled variable (sound intensity). Thus, you will observe all kinds of puzzling variability in the moth's dropping behavior that will likely be attributed to "error variance" or "random stimuli flying around" or whatever -- and you will miss the fact that the controlled variable is kept precisely at its reference level BECAUSE OF THESE APPARENTLY RANDOM VARIATIONS).

Obviously, I think you asked a very important question Avery. I hope this post helps a bit.

Regards   Rick

Date: Wed Mar 18, 1992 11:14 am PST
Subject: glossing over language differences

The perils of relying on English glosses to indicate meanings of morphemes in another language are often mentioned but the depth of our reliance on our native "background vernacular" is I think seldom really appreciated. Same principles apply in spades to invented metalanguages for semantic representation.

Here's an interesting case in point excerpted from the Linguist Digest.

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

6)
Date: Tue, 17 Mar 92 12:05:20 EST
From: cowan@uunet.UU.NET (John Cowan)
Subject: OVS

I have never really understood the necessity for talking of object-first languages, using this term as a cover for OVS, OSV, and VOS languages. What reason is there to believe that such a language actually has a different order rather than believing that it takes a different view of what its verbs mean? Using Okrand's study of Klingon as the readily-available example (:-)):

    puq legh yaS
    child sees officer
    The officer sees the child.

What reason is there to gloss "legh" as "sees" rather than "is-seen-by"? It seems to me a mere prejudice to believe that seeing is "inherently" more natural, and more deserving of a single morpheme, than being seen. So talk of the rarity of object-first languages can be reduced to talk of the rarity of "is-seen-by" as a single morpheme with "sees" as the derived form.

--
cowan@snark.thyrsus.com   ...!uunet!cbmvax!snark!cowan
                e'osai ko sarji la lojban

Linguist List: Vol-3-262.
<*>

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

This brings us rather sharply back to the old Whorf-Sapir(-Humboldt) hypothesis. Do differences in language correspond to differences in the world of perceptions in which we live and move and have our being? And the epistemological toughie: whatever your answer, how can we know that?

    Bruce
    bn@bbn.com

Date: Wed Mar 18, 1992 12:46 pm PST
Subject: Re: Mostly language

[From Bill Powers (920318.0800)]

Rick Marken: WHAT'S HAPPENING WITH THE SAN DIEGO MEETING???????

Bruce Nevin (920317) --

Emotion:

The term "emotional" isn't used much in any technical sense -- in science it's usually a pejorative meaning "you're letting considerations other than scientific truth into the argument." It also means "impulsive", or acting without higher-level reflection. In neither case do I think that anything is running open loop. Even the cockroach's escape "response" is an act that removes the danger. If it didn't it wouldn't have been learned or wouldn't have evolved.

Semantics:

One comment on your semantics remarks to Avery. The idea that words somehow refer to things "out there" is a SUBSTITUTE for the epistemology of control theory, in which the only knowable referents of words are perceptions IN HERE, which, of course are attributed to an objective external reality.

Correlations:

> ... correlation of an outcome with a purpose or goal smells >unscientific from a conventional perspective because the goal is a >perception hidden within the black box of the behaving system and the >controlled outcome is that subset of the possible sensory inputs to the >behaving system that are relevant to that goal.

True. But such opinions are based on a misunderstanding of science that assumes we mustn't guess about things we can't see. If that were how science really works, we wouldn't be able to explain why holding a magnet near a television set distorts the picture. If you guess carefully and quantitatively, you get a science like physics. If you guess sloppily and without high standards for accepting models, you get ... well, what you get.

[I'll return to "missing the point"]

>In your reiteration of basics, you say:

>>Ordinarily, we would expect action and disturbance to correlate
>>with the outcome.

>Do you mean the algebraic sum of action and disturbance correlates with
>the outcome, while neither correlates well, taken separately?

I meant that ordinarily we expect effects to correlate statistically with their causes. If there are multiple causes, like action and disturbance, we would expect each one to "contribute to the variance" of the effect. This is the bread and butter of statistical analysis. Conventional behavioral science is not prepared to recognize the case in which adding two causes together produces less effect than when either alone is used.

>What am I misunderstanding here?

Not much. I'm using the term "variance" in the statistical sense, not simply to mean "variation." In the statistical sense, variances don't add "algebraically." They add "in quadrature", meaning they add as the sums of the squares of the individual average variabilities. Statistical measures don't give significance to single data points.

In a tracking experiment, if you add the individual disturbance magnitudes to the corresponding action magnitudes point by point, you get the successive cursor positions. That's the algebraic sum you're talking about. To characterize the quality of control for an entire run, thousands of sets of data points, you can use statistical measures, like mean squared amplitude of the variables. The entire data set is used to obtain variances. You can compute the variance of the entire set of thousands of handle positions, the variance of all the disturbance magnitudes, and the variance of all the cursor positions. If there were no control, you would expect the sum of the handle and disturbance variances, the mass measures, to equal the cursor variance, another mass measure using all the cursor data points. When there is control, the actual cursor variance is much less than the expected variance -- the expected variance is 10 to 15 times the standard deviation of the actual cursor variance in a normal tracking experiment.

The only reason I have used these statistical concepts is to demonstrate to those who use statistics as the preferred way of analyzing data that something very strange is going on here. We've gotten a bit off the track by talking about how you can detect controlled variables when the reference level is randomly changing, using statistical methods like Marken's "mind-reading" demo. That's really doing things the hard way. While we may eventually have to do things the hard way, right now there are plenty of aspects of behavior we can study in which there aren't large unpredictable changes in reference signals -- especially if we design experiments to avoid them. In that kind of experiment, the statistical treatment is unnecessary, except to satisfy those who don't believe data unless they can see a correlation and a confidence level.

Missing the point:

When I use ill-advised expressions like that, I deserve the misunderstanding they generate. I think you and Avery know enormously more about control theory than any of your colleagues do, or for that matter than 99.99% of behavioral scientists do.

I was referring to something much narrower, namely, the point that I was trying to make about the difference between talking FROM the point of view of a theory and talking ABOUT the point of view. Bringing up the fact that your analysis and Avery's differ was an attempt to point toward a point of view superordinate to both. No takers yet. There's a hint of taking me up, however, in the following:

>Recourse to nonverbal perceptions in the perceptual hierarchy as the
>universe of "meanings" promises a way out of this circularity [of
>metalanguage]. The question remains: if the work can be done within
>language (with slight extensions to enhance regularity), why have
>recourse to a metalanguage purportedly outside the familiar language of
>words organized in sentences and texts?

I'm guessing that you're referring here, obliquely, to Avery's structural diagrams, which are not expressed in natural language. But can the work really be done within language? Doesn't it ALL have to be done within meaning? And isn't a structural diagram as good a way to refer to a meaning as a lineal sequence of words -- and in some cases, better? Diagrams in two dimensions would be best for referring to meanings in which more than one dimension of relationship exists at the same time. And lineal ordering is best for dealing with meanings in which time or sequence is a dimension. I don't know of any good way to refer explicitly to the structure of a program (other than just laying out the source code) except with a block diagram or some diagrammatic way of showing alternate paths at the same time.

Conventions and universals:

>There are some aspects of language that appear to be universal. There
>are other aspects that are defined by social convention. These
>conventional aspects of language are not universal.

The universal aspects, I presume, are universal because they arise from human properties common to all people, not statistically but without exception. If you have discovered universal aspects of language, then they are universal, period. Is this what linguistics has discovered? If so, there's no room for argument, is there?

>... having reached the desired "state where linguistic universals
>are trivially deduced from first principles," PCT still has to describe
>or account for those remaining aspects of language that are socially
>inherited and which do NOT follow from first principles.

PCT has no specific behavioral content, which is why it applies to all behavior. It doesn't tell you what variables people will control for, what perceptions they will become organized to experience, how stable their control will be, or what environments their actions will occur in. After you fill in the unknowns, it will tell you something about the consequences of your choices of variables and constants. So control theory has nothing to say about social conventions, linguistic or other.

<The means in the perceptual control hierarchy for learning, >maintaining, and orienting one's behavioral outputs to social >conventions or norms are presumably universal.

They are also the means for flouting, changing, and ignoring social conventions. Social conventions are descriptive, not prescriptive. They become prescriptive only when an individual accepts them as a means of achieving something, and what they prescribe is only what the individual takes them to prescribe. What I'm trying to emphasize is that the underlying machinery of the brain does more than language, and exactly what it does with language is optional (even though, through interactions, people tend to converge toward similar uses of language).

Expansions:

>[Expansion] uses at each point of word entry (operator entry) only the
>relatively small stock of words and word relations that are attested in
>many perfectly ordinary sentences of the language.

>It uses at each such point only reductions of word shape that are
>attested as minimal sentence-differences between pairs of perfectly

>ordinary sentences, such that the sentence-pairs informally meet my
>judgment of saying the same thing ...

This informal judgment brings in meaning: "the same thing" is "the same meaning." The terms, in addition to satisfying the mechanical or statistical word-relationship requirements, point to the same meaning as the reduced form. But if you can make this judgment, the implication is that the reduced form has the same meaning as the first expansion, the first has the same meaning as the second, and so on. Otherwise each expanded form would be adding meanings, not having the SAME meaning. This supports my scenario, not yours.

Yours:

```
       Heard (reduced) form  ----> expansion rules ----> Expanded form
              |                                               |
              |                                               |
          initial meaning  <-- The same? -->           expanded meaning
```

Mine:

```
       Heard (reduced) form                              Expanded form
              |                                               |
              |                                               |
          initial meaning  ---> imagination ---------> expanded meaning
```

So in your scenario, a rule is applied BY THE HEARER to the reduced verbal form, which has limited meaning, to produce an expanded verbal form, which has a more complete or explicit meaning. In mine, no such rule is applied BY THE HEARER.

What I'm suggesting is not that the transformation rules you propose are incorrect or unobservable, but that they may not be related to speech comprehension as you imagine them to be related. The situation I imagine could be represented this way:

```
                            Linguist-->   Expansion rules
                           / \
                         /     \
                       /         \
                     /             \
                   /                 -------------------
                 /                                       \
               /                                           \
       Heard (reduced) form                              Expanded form
              |                                               |
              |                                               |
          initial meaning  ---> imagination ---------> expanded meaning
```

The linguist can certainly derive objectively defensible expansion rules which, applied to the reduced verbal form, will yield expanded verbal forms. But the observer can only imagine that the hearer is actually applying those rules. If the hearer is doing the expansion as I suggest, then the rules seen by the observer analogize but do not describe what is going on perceptually in the imagination of the hearer. The processes I label "imagination" explain why the reduced form and the expanded form are related as they are observed to be related. They APPEAR to be related by a rule in which words, not perceptions, are the arguments. But in fact, in the hearer, words are not involved in the expansion until it is finished. I say "in fact," but of course this is just a proposal.

In order to distinguish your scenario from mine, and show that the hearer is actually applying the rules, it's necessary to do more than show that the rules fit what is observed. It's possible to devise a set of rules that will describe any process in words or symbols, even if that process uses neither words nor symbols. Every computer simulation of a physical process illustrates this fact. Nobody believes that the computing steps are actually employed in the physical process. They are simply equivalent in their final effects to the processes actually at work.

I'm not picking on you: the same alternative applies to Avery's diagrams. By seeing these approaches to representing language structure as analogies, I think one can get a better sense of the machinery underlying language phenomena (and behavioral phenomena in general). There must be a relationship between the linguistic analyses and that underlying machinery. But I doubt very much that it's a one-to-one correspondence -- that linguistic processes are specifically about words, or that they are actually carried out in words.

I'm not trying to put down anyone's approach -- just to get a higher-level look at what's going on.

Best,   Bill P.

Date: Wed Mar 18, 1992 2:16 pm PST
Subject: Re: off the track

[Martin Taylor 920318 16:30]
(Bill Powers 920318 08:00)
>
> We've gotten a bit off the track
>by talking about how you can detect controlled variables when the reference
>level is randomly changing, using statistical methods like Marken's "mind-
>reading" demo. That's really doing things the hard way. While we may
>eventually have to do things the hard way, right now there are plenty of
>aspects of behavior we can study in which there aren't large unpredictable
>changes in reference signals -- especially if we design experiments to
>avoid them. In that kind of experiment, the statistical treatment is
>unnecessary, except to satisfy those who don't believe data unless they can
>see a correlation and a confidence level.

By doing experiments with fixed reference levels, I think you are leaving yourself wide open to an S-R interpretation. If in an ECS the reference level is fixed, then the "stimulus" percept does determine the error signal and thus the observed behaviour. That there are disturbances to be compensated for is clear, but interpretable as behaviour that is a consequence of the disturbing stimulus.

Separate comment on the same paragraph:

Statistical treatment is avoidable only if you have almost correctly guessed what percept is being controlled and, perhaps more important, what effects your experimentally induced disturbances would have if the subject were not controlling. If you are in error in either of these, you are in trouble without statistics. But I agree with you that "confidence levels" are an abomination, though my reasons are quite different.

And I don't think you should believe that psychologists are unaware of the effects of correlation (negative or positive) on the variance of joint phenomena.

To change the subject: is no-one interested in my argument that the degrees of freedom problem leads to modular reorganization, or did it not get out to the list?

Also, no-one has made any suggestions in respect of my Paris talk, which was sent at around the same time, which makes me think that perhaps neither item got distributed (and perhaps others as well). I would think Paris is an opportunity to put the PCT case before a reasonably influential audience, and you would like the opportunity to suggest lines of attack.

Martin

Date: Wed Mar 18, 1992 2:42 pm PST
Subject: Degrees of Freedom

[from Gary Cziko 920318]

Martin Taylor 920318 16:30 says:

>To change the subject: is no-one interested in my argument that the degrees
>of freedom problem leads to modular reorganization

I just got a phone call from Bill Cunningham in Fort Monroe, VA, who said he is very much interested in this topic but due to computer reorganization is unable to respond to CSGnet, although he does receive mail.

He wanted me to let people know it may be a few weeks before he can communicate with CSGnet and hopes that this topic stays alive.--Gary

Date: Wed Mar 18, 1992 3:58 pm PST
Subject: Re: off the track

[From Rick Marken (920318b)]

Martin Taylor (920318 16:30) says:

>By doing experiments with fixed reference levels, I think you are leaving
>yourself wide open to an S-R interpretation.

I agree, actually. If one's goal is to convince people that they are missing something important (the secularly adjustable reference for perceptual inputs) by imposing an SR interpretation on control phenomena, then its nice if you can show that there is a varying reference and that it is determined by the subject. I am working on an experiment now that

I hope will demonstrate this phenomenon; it is a computer demo; my experience is that it is quite hard to develop multilevel computer demos of control. I still like Bill's "portable demonstrator" the best.

>To change the subject: is no-one interested in my argument that the degrees
>of freedom problem leads to modular reorganization, or did it not get out
>to the list?

I am very intereted in what I think of as the "degrees of freedom" problem. But I think we have different ideas about what the degrees of freedom problem might be. Maybe we could try again -- what is the degrees of freedom problem? What is modular reorganization?

>Also, no-one has made any suggestions in respect of my Paris talk, which was
>sent at around the same time, which makes me think that perhaps neither item
>got distributed (and perhaps others as well).

I don't remember whether it was distributed or not. But I imagine that, if it was, my suggestion then would have been the same as it is now -- TAKE ME!!!!

Regards   Rick

Date: Wed Mar 18, 1992 9:36 pm PST
Subject: repeats

[Martin Taylor 920319 00:30]

We had mailer problems, apparently, last week and particularly last weekend.

I am going to re-send four messages, two from March 11 and two from March 16. The two from March 16 apparently did not get out to CSG, because I asked about them explicitly. The others may have done, but I don't remember getting any responses about them. If you have seen them before, please forgive the repetition.

Martin   - One from 3/11 deleted as redundant. Dag

Date: Wed Mar 18, 1992 9:54 pm PST   (Originally March 11)
From: Control Systems Group Network
Subject: Low correlations

[Martin Taylor 920311 11:15]
(Bill Powers 920310 17:00)
>
>Martin Taylor (920309), Bruce Nevin (920310) --
>
>You guys are going all around the point here. I have a distinct feeling
>that you're avoiding it. On the other hand, there may be some critical fact
>that I don't seem to be communicating, so let's try again:
>

I'm not sure which way the lack of communication goes, but it sure seems to be there. I see nothing in your posting, whether by diagram or in the explanations that is not crystal clear, and has been so for ages, at least in my mind, if not in my writings. It's the starting point for most of what I have been trying to write over the last couple of months. What might be the point that Bruce and I are missing? Could you put it another way?

I think that if you take your posting together with the following paragraph from mine of 920309 11:00, you may see why I think we have a communication problem, and perhaps will be able to resolve it:

>If, as an experimenter, one can presume some pattern in the mutually observable
>environment represents a perceptual variable being controlled by the subject,
>then one can attempt to disturb that pattern and see whether the subject
>acts so that the pattern is restored or maintained. The pattern will show
>little correlation with the experimenter's disturbances or with what the
>experimenter observes of the subject's actions. If the experimenter happened
>to be correct that what she did would have disturbed the pattern if the
>subject had not been there, then there is evidence that the subject is
>controlling. The presumption that the experimenter would have disturbed the
>pattern is just that, a presumption. It is not an observation, because it
>didn't happen. Explaining why things do not happen is trickier than providing
>rationales for why they do happen. The failure of a presumed "cause" is
>easier to justify as that it was not a cause than as that an exactly
>countervailing cause was applied at the same time. I think this is at the root
>of the communication difficulty with cause-effect psychologists. Causes have
>effects, and PCT is supported when what should be causes are observed to
>have little or no effects.

Note the words in the 6th and 7th lines: "If the experimenter happened
>to be correct".

At the same time, I would love to pursue my degrees-of-freedom discussion, but I can't until the question of zeros is resolved. Could you explain the matter of the error signal for sequence, which is one part of the remaining problem? Rick has partly resolved the other part--that the spreadsheet provides a counter-example--by pointing out the considerable non-orthogonality among the ECSs in the spreadsheed. The necessary non-zeroing related to non-orthogonality was to be part of the later discussion, but I had forgotten that it could come up even when the degrees of freedom for input and output of the hierarchic CS are in balance.

Martin

Date: Wed Mar 18, 1992 9:55 pm PST
Subject: Re: motor programs etc.

From mmt Mon Mar 16 11:28:52 1992

[Martin Taylor 920316 11:10]
(Rick Marken 920312)

I seem to be reconnected again, but over the weekend there has been only this one item from Rick and another from Bill. Is this right?

Rick says: I think fair is fair --
>conventional psychology is pretty content to ignore PCT; I think PCT
>should just return the favor.
>

One of the few things I was told in graduate school that has stuck with me as something I still believe: If there are two schools of thought, each with good reason claiming that they have the truth and the other doesn't, they are probably both right, except in that claim.

It is obvious to me that the basic ideas of PCT have to be right, just as are those of (say) signal detection theory or information theory. PCT cannot work if the informational requirements are not met. The actions that a subject performs in a signal-detection study cannot be accounted for without PCT.

For some weeks I have been trying to get this question of zero references in a stabilized hierarchy sorted out, so that I can get to my main point--that most perceptual activity at any moment in time is passive and uncontrolled. PCT has to acknowledge this--it is required by the informational arguments and can't be wished away by faith or dogma.

Quite apart from that, it has to be the case that consistencies observed in conventional psychological experiments tell us something about what goes on inside a person, because consistencies represent something that resists whatever (unknown) disturbances to which the subject is subject (!). It may be, and usually is, hard to know what these consistencies tell us, especially if we don't look at what the subject is trying to achieve (i.e. if we ignore the PCT approach of looking at controlled perception). But the difficulty of discovering what a poorly conceived experiment says should not be sufficient reason for denying that it says anything useful. It may be easier to regather the information in a better conceived experiment, but as amply testified by the discussions on this net, designing such experiments in PCT terms is not easy, either.

There's plenty of bath water, but there are babies in it, so be careful what you throw away.

Martin

Date: Wed Mar 18, 1992 10:01 pm PST
Subject: Assistance requested for talk

From mmt Mon Mar 16 11:39:16 1992

[Martin Taylor 920316 11:30]

This is a generalized request for assistance that applies until mid-August.

I have been asked to give a 3-hour keynote address to a high-level "Autumn Summer School" on human-computer interaction in Paris in early September. Naturally, I intend to take a PCT position as a base for looking at all the different approaches to HCI that might come up, but this will probably be a very general base, rather than dealing with specifics of the theory. I mean that I will probably emphasize throughout that the user is doing whatever actions he or she performs in order to be able to perceive some desired thing. The computer output exists in order that the appropriate things can be perceived and compared with the results of imagination as well as with the reference signals.

My request is for CSG people to suggest themes that I might address, or issues in HCI that have posed problems because they have not been addressed from a PCT view, or, if it is feasible, more detailed discussions of aspects of HCI that have been or might be addressed by PCT.

This may be a good opportunity to get PCT thinking into a mainstream community of applied psychology. Then again, it may not, but it's worth trying, I think.

Martin

Date: Wed Mar 18, 1992 10:01 pm PST
Subject: GAs, modularity of reorganization
From mmt Tue Mar 17 11:23:40 1992
Subject: Re: Alpha control; Genetic Algorithm; BEERBUGS

[Martin Taylor 920317 10:45]
(Bill Powers 920316.1600)

>A sixty-millisecond
>action doesn't seem too short for feedback control. Even the human arm
>positioning system can move the forearm from one stable position to another
>in around 100 milliseconds under active control (although that's the lower
>limit and involves tremendous muscle forces -- 20g accelerations). In a
>cockroach where the path lengths are only a few percent of those in the
>human spinal systems, I should think the loop delays can't be more than 5
>milliseconds, which allows plenty of leeway for performing a movement in 60
>milliseconds under good feedback control.
>

A long time ago, I was told that the fastest reflex in the human was the eyeblink "response" to a puff of air: 5 msec. Interpreted as Bill interprets the cockroach avoidance turn, the control is for not feeling the air (or a touch?) on the eyeball.

A good pianist can play smooth sequences of notes at around 60 msec intervals. There may be predictive control operating here--it is hard to stop playing a figure that has been started--but (in contrast to lest year) I see no evidence that the individual finger strokes are not controlled.

On the random-walk of reorganization: there is a real degrees-of-freedom problem here. In the 2-D demo, there is a good probability that the walk is within (say) 60 degrees of the direction to the target. In 1-D, the probability is 0.5, in 2-D 0.33, and goes to zero for very large dimensionality. To exercise the control, you have to get the point to move again if it goes the wrong way, which is 50% of the time whatever the dimensionality. But in large dimensionality it hardly ever goes very close to the direction you want, which means that the control is very slow. Almost always, the random move is very nearly orthogonal to the direction you want. My interpretation of this is that reorganization can work well only if it applies to a small number of connections at a time--for example within a small modular group of ECSs, or coordinated connections that link a small number of modules that are themselves unaffected by the reorganization.

According to this view, the same problems that affect GAs also affect reorganization, and the solution is the same. GAs modularize by sequence inversions that lead to co-location of cooperating genes; the hierarchic control system modularizes by reorganizing small groups of ECSs so that they can perform specified functions in a coordinated way, and then modifies only the links among those groups.

The same solution seems to have been used by natural evolution--we did not spring full-formed from the brow of Zeus, but were built from conspiracies of successful sub-organizations.

Martin

Date: Wed Mar 18, 1992 10:55 pm PST
Subject: plummetting moths, drooling dogs

To continue the discussion of moths & S-R systems, what would be the story about a dog that was hardwired to salivate when it smelled food? Abstractly, these two systems look alike w.r.t. what is inside the organism:

```
       _______          _______
   --->| det- |  Y/N    | eff- |  --->
   --->| ect- |-------> | ect- |  --->
   --->| or   |         | or   |  --->
        ------           ------
```

e.g. a vector of sensory inputs gets transduced to an approximately ON/OFF signal, the ON value of which sets off an effector. The difference is that in the moth case, the presumed result of the effector is to make the sonar signal go away, while in the salivation case, the effector does NOT make the smell of food go away. But as far as what's inside the critter itself, the setups look pretty much the same.

And furthermore, is it so obvious that the moth's system is actually making the signal go away? What it might be doing is making the moth undetectable by getting it on the ground as fast as possible (I don't know whether the dropping actually normally gets the moth out of sonar range or not).

So it seems to me that exactly the same hunk of circuitry might be appropriately regarded as an ECS or an S-R system, depending on whether there is an R-S connection that is an important aspect of the functioning of the system. E.g. if the dropping systems serves get the moth out of sonar range fast, it's an CS, if it gets it to be undetectable against the ground, though still in range, it's S-R.

Avery.Andrews@anu.edu.au

Date: Thu Mar 19, 1992 8:33 am PST
Subject: an unusual event

To: CSGnet members in general
Subject: an unusual event
Date: 03/19/91
From: David Goldstein

Yesterday, at the dentist's office, my dentist told me a strange story. It concerned his six year old son. This well behavied child, one day, pulled the fire alarm. The firemen came. When the child saw the commotion, he told his teacher that maybe he was responsible. The teacher asked him why he pulled the fire alarm. He said: I don't know. I wasn't thinking about it before I did it. It turned out that the fireman discovered a fire was in process in a locked closet containing paints and stuff like that. They said that an explosion would have occurred shortly. The school is a Catholic one and they are putting it off on a miracle. The discussion of the moth which has been going on lately seems related to this unusual event. Any body care to speculate on the controlled variable involved?

Date: Thu Mar 19, 1992 9:42 am PST
Subject: Re: an unusual event

Experience suggests to me that there are forms and sources of information and ways of perceiving it that are not represented in our conventional models of what is going on (consensual reality). This is not to say that alternative models have greater claim to Truth, but only that it is very likely that a great deal of our perceptions get ignored just because higher-level input functions are indifferent to them.

Date: Thu Mar 19, 1992 11:29 am PST
Subject: unusual event

This unusual event of the dentist's son is interesting for a number of reasons. It's not difficult to understand how it happened in a general manner--something is perceived, an error signal is created, an action is taken--all through various levels of the hierarchy. That's easy. And we certainly agree that we are not aware of the reference levels at each level of a wholistic behavior--getting a caffiene buzz vs. lifting a coffee cup. So the fact that the boy didn't know why he hit the fire alarm is not in itself so unusual. But I have a very undeveloped notion of conscious awareness and error and automaticity which I've developed from the Vallacher and Wegner article, and this notion would have me expect that the boy would have awarness of "why." So the consciousness question is intriguing.

Bruce addressed the perceiving end of this event--that our present models don't account for the possible varieties of perceiving.   Now it may be the case that the boy sensed by smell or touch that something was amiss and the whole process models as normal, only the perception is not conscious. Or is it that there is a different perceiving mechanism or organ that we don't know about? There are certainly some things that we often say demonstrate "ESP" or the like which can be explained by normal perceiving mechanisms where consciousness awareness of the "stimuli" is not present (perceiving that someone is standing behind across the room by means of underdeveloped echolocator skills). But there are some things which don't fit under such explanations, many which I have experienced first hand to my amazement. I have learned some of these (intuitive) skills and can do such fun things as finding someone's birthday by means of holding and watching a pendulum. Now I know how utterly ridiculous that sounds but I say it anyway cause I know what my experience. Such acts, unfortuneately, cannot be done when the internal environment is such that we define it as being "critical, analytical." Learning such skills required that one learn how to not analyize for at least a while (a difficult task for myself, probably less so for most). Now I realize that my terminology is extrememly vague and my statements seemingly unfalsifiable, and for this I apologize. But there's something here to be examined, if it is examinable.

So I wouldn't call it a miracle. Any comments?

Mark Olson
m-olson@uiuc.edu

Date: Thu Mar 19, 1992 11:43 am PST
Subject: an unusual event

>From: David Goldstein
>
>well behavied child, one day, pulled the fire alarm. The
>firemen came. When the child saw the commotion, he told his
>teacher that maybe he was responsible. The teacher asked him
>why he pulled the fire alarm. He said: I don't know. I wasn't
>thinking about it before I did it. It turned out that the fireman
>discovered a fire was in process in a locked closet containing
>paints and stuff like that.

How did the firemen stumble on the fire? Perhaps the same clues?
I should think that a subliminal smell of fire could suffice.

Bill
--
Bill Silvert at the Bedford Institute of Oceanography
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2
InterNet Address: bill@biome.bio.ns.ca

Date: Thu Mar 19, 1992 3:08 pm PST
Subject: Misc responses

[From Rick Marken (920319)]

Martin Taylor (recently) says:

> If there are two schools of thought, each with
>good reason claiming that they have the truth and the other doesn't, they
>are probably both right, except in that claim.

Well, I never claimed to have the truth. I claim that HPCT is the best current explanation of a phenomenon that is currently not studied in psychology -- except obliquely -- control. So I claim that that model is better than any other as an explanation of that phenomenon. I guess I am also saying that most of what psychologists think of as the phenomenon called "behavior" is actually control -- so I am saying that my model is the correct model of behavior (as far as we can go towards correctness with current know- ledge) and their models are, thus, wrong.

>It is obvious to me that the basic ideas of PCT have to be right, just as
>are those of (say) signal detection theory or information theory.

I think I don't agree -- maybe I don't understand.

> my main point--that
>most perceptual activity at any moment in time is passive and uncontrolled.
>PCT has to acknowledge this

I don't know about "most perceptual activity at any moment in time" being passive but has any PCT control theorist ever denied that a great deal of our perceptual experience is not controlled? I certainly don't deny it -- I'm perceiving all kinds of variables that I cannot influence in any way except (if they are visual or auditory) to make them go way (closing eyes, plugging ears). There are many that I'm not controlling that I could control (eg, the color of my walls) and others that I will never be able to control (the number of clouds in the sky). I love those uncontrolled perceptions.

>   --it is required by the informational arguments
>and can't be wished away by faith or dogma.

You got me here? What are the informational arguments? And why should they matter to me? (I sound like the disobediant child at passover).

>There's plenty of bath water, but there are babies in it, so be careful what
>you throw away.

I agree -- but I think it's more like dimes in the great salt lake. We do try to keep an eye out for 'em (those of us who don't mind opening our eyes in salt water; I don't mind doing it but I'd rather look for my dimes at the bank).

>I have been asked to give a 3-hour keynote address to a high-level "Autumn
>Summer School" on human-computer interaction in Paris in early September.

You luck!

>Naturally, I intend to take a PCT position as a base for looking at all the
>different approaches to HCI that might come up, but this will probably be
>a very general base, rather than dealing with specifics of the theory. I
>mean that I will probably emphasize throughout that the user is doing whatever
>actions he or she performs in order to be able to perceive some desired
>thing. The computer output exists in order that the appropriate things
>can be perceived and compared with the results of imagination as well as
>with the reference signals.

Excellent! Don't forget to mention those disturbances; mistypings, different results depending on context, etc!

>My request is for CSG people to suggest themes that I might address, or
>issues in HCI that have posed problems because they have not been addressed
>from a PCT view, or, if it is feasible, more detailed discussions of aspects
>of HCI that have been or might be addressed by PCT.

Well, it's not very detailed but I did write a little article for the Human Factors Society Bulletin (the December 1986 issue, I believe) that was about control theory and HCI. I thought it was great. Some others did to -- but the conventional types spotted the heresy immediately and suggested that I get with the "real" research program -- looking for the effects of input variables on responses.

If you can't get a hold of a copy (call the HF Society in Santa Monica,CA maybe) I can send you a reprint if you're interested.

>On the random-walk of reorganization: there is a real degrees-of-freedom
>problem here. In the 2-D demo, there is a good probability that the walk
>is within (say) 60 degrees of the direction to the target. In 1-D, the
>probability is 0.5, in 2-D 0.33, and goes to zero for very large
> dimensionality.

Actually, we have never looked at efficiency as a function of the dimension- ality of the space in which the value of the one dimensional controlled is computed. Might be interesting. Maybe Bill Powers has some data on that?

Avery Andrews (920319) says:

>To continue the discussion of moths & S-R systems, what would be
>the story about a dog that was hardwired to salivate when it
>smelled food?

> in the moth case, the presumed result of the effector is to
>make the sonar signal go away, while in the salivation case, the effector
>does NOT make the smell of food go away. But as far as what's inside
>the critter itself, the setups look pretty much the same.

Well, you are assuming your conclusion. If smell caused salivation as you postulate and salivation has no effect on smell, then we've got an SR system and it doesn't look like anything is controlled. I'm not sure it actually works that way, though. First, I'm not sure that there is a hard wired connection between smell and salivation. I believe that the "natural stimulus" for salivation is stuff in the mouth that tends to absorb liquid. So I'd say sensed "wetness" is the "stimulus" for salivation -- that is my assumption; I believe that you might test this by seeing if a dog salivates to smell when it's mouth is full of water or if it salivates when an orderless dry substance is placed in the mouth. I predict no to the first and yes to the second.

If smell is involved in salivation (without training) then I would have to be convinced that salivation really has no effect on the smell signals -- it's plausible that it does have such an effect though; salivation probabaly deceases the amount of "smell chemicals" that go back up into the nasal passages from food in the mouth.

>And furthermore, is it so obvious that the moth's system is actually
>making the signal go away?

No. But falling doesn't have to make the signal go away -- it just has to affect it -- which it obviously does.

> What it might be doing is making the moth
>undetectable by getting it on the ground as fast as possible

Of course it DOES. But that's not what the moth is doing. That is an Irrelevant Side effect (ISE -- there, got it right this time) that happens to be VERY relevant to the moth's survival. The moth just knows nothing about it. If the falling put the moth into the hand of a friendly bug researcher (like Bill Powers and his pet beetle) that would be fine with the moth too -- just get that damn signal outta here, sayeth the moth's control system.

>So it seems to me that exactly the same hunk of circuitry might be
>appropriately regarded as an ECS or an S-R system, depending on
>whether there is an R-S connection that is an important aspect of
>the functioning of the system.

You betcha!! Now think about how much you can "do" without influencing one or another of your sensory inputs. When you realize that that amount is almost precisely ZERO then you have entered the PCT zone.

Regards Rick

Date: Thu Mar 19, 1992 4:24 pm PST
Subject: Taylor catchup

[From Bill Powers (920319.0900)]

Martin Taylor (920318) --

Re: coin game

So there's another wheel patent down the drain. Did Garner use this as a method for discovering controlled variables, or as an illustration of the problems that arise in carrying out such explorations? I'd appreciate it if someone with access to a bigger library would look up the Garner reference and say something about it briefly from the CT point of view.

>If, as an experimenter, one can presume some pattern in the mutually
>observable environment represents a perceptual variable being >controlled by the subject, then one can attempt to disturb that pattern >and see whether the subject acts so that the pattern is restored or >maintained. The pattern will show little correlation with the >experimenter's disturbances or with what the >experimenter observes of >the subject's actions. ...

The Test usually isn't done in such an arm's length way. Usually you pick a potential controlled variable because you can see that physically the subject's actions ought to be having an effect on it, and you can also see that disturbances can have an effect on it. You also have your own experience to draw upon for starting guesses: if I were acting like that, what would I be controlling for? You're right in saying that it's necessary to know how much effect a disturbance ought to have on a variable if there's no control. Usually, however, the difference between control and no control is so large that a ballpark estimate or just previous experience with that kind of variable is good enough.

>The presumption that the experimenter would have disturbed the
>pattern is just that, a presumption. It is not an observation, because

>it didn't happen.

Again, too abstract an approach. When you disturb a coin, it will stay disturbed until the subject corrects the error, if any. For faster control systems, when you see that the variable doesn't change, you also see that the subject's action on it DOES change. You have more evidence to go on than just the failure of the variable to change. Even without predicting how much the variable should or might have changed, you can (often) block the subject's ability to perceive the variable, and find that now it changes. So beside just the failure of the variable to change, you have information from relationships between the subject's actions and the variable, and between the subject's perceptions and the variable. The Test incorporates all these factors, as presented in BCP.

I did reply about the "sequence error" problem; the gist of my reply was that I don't know how to design a realistic sequence control system. The basic requirement is that the perceptual function provide a signal that indicate a certain sequence in progress, the signal perhaps growing as more and more correct elements appear and declining when an incorrect element appears. This is an experiential requirement: when I start to spell M-I-S-S-I-S-S ... you have a pretty good idea what the sequence is long before it's finished. Also, when there's a repeating sequence like tick-tock-tick-tock --- you get a sense of the same sequence being present as long as it continues, so we need a steady signal while the sequence is in progress. An extra or missing tick provides a brief error signal. Anyway, I don't have a lot to say about sequence control -- just that it happens.

If a controlled sequence is in progress, it is maintained in progress by small differences between the produce-this-sequence reference signal and the this-sequence-in-progress signal. That's the best I can do.

>If there are two schools of thought, each with good reason claiming >that they have the truth and the other doesn't, they are probably both >right, except in that claim.

PCT has to be compared with other points of view at the same level. That is, PCT should be compared with S-R theory or field theory or cognitive theory, not with such things as psychophysics or neural network models. When you get below the level of overall organization and start looking at how the components work (so their overall organization doesn't matter), you're asking how the components of either a control system or an S-R system work; there's no difference at that level. Signal-detection theory is about the perceptual systems, and control theory doesn't force us to accept any particular model of perception -- just whatever one is best. If information theory can tell us something about bandwidths and siognal-to-noise levels and probabilities as they appear anywhere in the model, fine. That won't change the model's organization, which is what matters.

PCT and S-R theory do have a link. It's possible to show that the "stimuli" of S-R theory, in most but not all cases, are better thought of as disturbances, so we will at least look for controlled variables being stabilized by the "response." S-R theory is then predicted by PCT, as the relationship between a disturbance and an action.

But nobody likes to be told that his or her life's work is a special case of someone else's theory.

>For some weeks I have been trying to get this question of zero >references in a stabilized hierarchy sorted out, so that I can get to >my main point--that most perceptual activity at any moment in time is >passive and uncontrolled.

From what I remember of your previous remarks about zero references, the problem seems to come up because of thinking in digital rather than analog terms. In digital terms, a high-level variable is there or not there, so the error is either present or absent. If you look at any particular example of such cases, you can see how the apparently digital variable can be subject to analog disturbances. Elements of the perception aren't just right or wrong: they can be almost right, or a little wrong. These differences call for variations in lower-level reference signals to keep them from getting large enough to constitute a serious error. Maybe this is where fuzzy logic should come into the model, to get away from these either-or concepts that cause conceptual problems. The main problem they cause is this: if there's an error, an action is started that corrects the error. But if it corrects the error, the action will stop, which causes an error again, and an action, and no action, and an action ... Beginning servomechanism engineers often start out this way, trying to understand control systems qualitatively, with the result that they can't see how the system could ever find a stable state.

When you see all variables as continuous, even logical ones, you can now have error signals of different sizes, and equilibrium becomes possible. The equilibium occurs not at zero error, but at a very small amount of error which, as it fluctuates, produces the adjustments that keep the error small.

As to the second point, I think I've agreed with it before. Most perceptions are not controlled; even among the controllable ones, not all are being controlled at one time. I think I used the example of controlling arm position: all you really need to control is elbow and wrist position, in fact one point on elbow and wrist, to determine the arm's configuration in two degrees of freedom. But you can still see all the points on the arm in between. External constraints force all those intermediate points to change as the elbow and wrist positions change. You could imagine a very elaborate control system that required every point on the arm as perceived to match a corresponding point on a reference-arm, but this would be enormously wasteful redundancy. It isn't necessary to have a reference signal and a perception for every point on the arm.

The same is probably true of all controlled variables. What is controlled is only what is necessary to control. Perhaps in a more advanced model we might want to allow one level of control to select control points among the variables of lower level, different control points being selected even for the same (global) controlled variable, depending on what other control systems are acting at the same time. In some circumstances you might want to control just your hand, letting the elbow go wherever it wants to, while in another circumstance -- holding a newspaper under your arm -- you'd want to pick different control points to constain where the elbow is. I don't think we're ready for a model that elaborate, however -- maybe two generations from now.

Keynote address:

Three hours! I suggest strongly that you get a projection plate and show them some demos on a big screen.

Possibly another subject of interest might be writing error-free programs. I've always thought that there isn't much distance between current programming practices and an HPCT approach. Instead of treating the computer as an open-loop device, monitor every intermediate result and compare it with a reference signal to make sure it's of the right kind, makes sense in some terms, and so on. Of course this doesn't mean comparing the result of computing 2 + 2 with a reference signal of 4, but something more subtle, like checking that the sign of the result is consistent with the signs of the arguments, and so on. A lot of this is done already -- overflow-checking, range-checking and so on --but conceiving of the process as one of controlling for critical perceptions instead of just commanding things to be done might lead to some new and more reliable programming methods.

An incident you might want to cite is the 3-mile-island accident, which arose in part because one indicator showed the status of a command signal instead of the status of the result (flow of water through a valve). Wrong perception under control. The indicator said valve closed, but the valve was open, letting the cooling water out (but check that, I'm not sure which way the error was). The principle is, you can't control what you can't perceive. I think organisms are so full of feedback connections because basically you can't trust nature. If I tell my finger to move, I want to SEE it move and FEEL it move. Then maybe I'll believe it really moved. This principle, while it seems very suspicious and fussy, seems to have resulted in a remarkably competent mechanism.

Ah -- a pet peeve about instructions for using computer programs. A lot of programmers will prepare a handy list of what all the keystrokes do, but the list is ordered the wrong way. Down the left side of the page you have control-a, control-b, ... control z, F1, F2..Fn, and so on, in nice neat keyboard order. So if you want to know what a given keystroke does, you can quickly find the key and look up its action.

But if you want to know what output to produce to create a preselected result, you may have to read every entry on the list. PCT says that we have reference levels for results, not for the actions that produce them. We start by wanting to begin a block define, not by wanting to press F4. So these handy lists should be organized by what is to be accomplished, not by the action that achieves the result.

If anything else pops up I'll let you know. Actually, I would think that getting everything you could think of yourself into only three hours would be a real challenge!

Random Walk:

>On the random-walk of reorganization: there is a real degrees-of->freedom problem here. In the 2-D demo, there is a good probability >that the walk is within (say) 60 degrees of the direction to the >target. In 1-D, the probability is 0.5, in 2-D 0.33, and goes to zero >for very large dimensionality.

Brilliant. You're right. I've always had the feeling that we can't get away with just one global reorganizing system, because what would make it randomly reorganize the right thing? My attempt to deal with this was the postulate about attention being drawn to error, and the locus of reorganization following attention. But that can't handle aspects of the system that aren't available to awareness, such as the damping coefficient in a limb control system.

I think your conclusion about the required modularity of reorganization is correct.

Actually, E. coli steers just fine in three dimensions, so that number of degrees of freedom isn't a problem. But as you introduce more and more of them, the selection criteria become a REAL problem unless you have independently applied selectors operating on different dimensions of variation. In the limit you could have one reorganizer per control system -- or more. E. coli, by the way, can chemotax toward or away from something like 27 substances, using only the one random-tumbling output. So the key is clearly in the perceptual selection process, not in the output process.

In a vague way I've realized that the environment and the basic behavioral machinery have to have some special properties for reorganization to work. The least is that small reorganizations must have small effects. In E. coli, if the next direction of movement is onlyh slightly different from the previous one, the time-rate-of-change of concentration that's sensed must differ only slightly from the previous one. The geometry of space and the properties of diffusion see to it that this is true. How can we translate this into a general requirement? This will tell us something about the initial organization of the nervous system that evolution has to provide if reorganization is to be possible. Does this sound like your kind of problem?

Date: Thu Mar 19, 1992 6:34 pm PST
Subject: little stick man

[From Wayne Hershberger]

(Bill Powers 920312)
>There was a big breakthrough this week [re: little stick man].
>To explain it, I have to describe the basic spinal-cord part of
>the model. I won't go into every detail.

>Pulling on the load point will cause the muscle to shift the
>load point in the direction of the pull (one muscle will pull in
>the direction of the load while the other relaxes). As a result,
>the sensed force in the tendon will remain the same. This system
>reduces the effective mass of the arm nearly to zero. It will
>apply the specified force regardless of the position of the
>load.

Congratulations on your breakthrough! This is VERY important work. Could you please expand upon your explanation? I am not sure that I understand your use of the terms _load point_, nor your expressions _pulling on the load point_ and _to shift the load point_. Nor do I understand your parenthetical expression above; which muscle relaxes?.

What values are you using as reference inputs for the gama system?

What is the magnitude of the error signal in the position loop when the arm comes to a stop ; does the magnitude of this error signal vary systematically with the eccentricity of the arm's position?

Warm regards, Wayne

Wayne A. Hershberger   Work: (815) 753-7097
Professor of Psychology
Department of Psychology   Home: (815) 758-3747
Northern Illinois University
DeKalb IL 60115   Bitnet: tj0wah1@niu

Date: Thu Mar 19, 1992 8:59 pm PST
Subject: plummetting moths; drooling dogs

[from Avery Andrews (920320)]

(Rick Marken 920319)

>just get that damn signal outta here, sayeth the moth's
>control system.

Well, that's exactly what I'm conjecturing that maybe is *not* happening, because the changes in the signal induced by the moths behavior aren't material to survivability of the moth. So my judgement so far would be that if it what is going on is avoidance of being eaten via quick departure from the vicinity of the bat and the accompanying sonar signal, its control, otherwise it ain't. I'm not committed to its actually being one way or another - just to clarifying the issues.

You might, of course, be right about the salivation stuff - If it works the way you say, then I entirely agree that that is control. Suppose salivation in response to smells is acquired. Then that might be part of a control system designed to forestall error signals, on the general basis that smell-of-nice-garbage is a harbinger of dryness-in-the-mouth, which gets forstalled by some anticipatory salivation. I am intrigued that this looks like control on a larger time scale, an S-R hookup on a smaller one. The case of hardwired salivation-in-response to smell doesn't look control-ish except perhaps at the level of evolution.

>You betcha!! Now think about how much you can "do" without influencing
>one or another of your sensory inputs. When you realize that that amount
>is almost precisely ZERO then you have entered the PCT zone.

But when these influences have no material relevance for what is being done, I wouldn't want it to call it control. I'd agree that the role of S-R hookups in human psychology is pretty minimal, but that does not equal necessarily nonexistent (I won't repeat my previous suggestions about possible candidates). *I* think it's should be quite useful to look at cases of possible S-R hookups with an open mind, & describe the kinds of facts that would induce one to classify them as one or the other (as I think we're making progress on). I see no reason why simple critters shouldn't have a fair number of these things rattling around in their circuitry.

Avery.Andrews@anu.edu.au

Date: Fri Mar 20, 1992 6:19 am PST
Subject: replies to Bill, Avery

[From: Bruce Nevin (Fri 920320 08:41:01)]

(Bill Powers 920318.0800) --

By emotion I didn't mean "emotional" as opposed to "logical" but rather was recalling a survey article in _New Scientist_ that I had described to the list a few months ago. I believe the researchers supposed that the "responses" were open loop, and from a CT perspective this has to be wrong. More interesting is the idea of stereotyped categorization of experiences leading to very quick setting of a goal plus neuropeptides like adrenalin cranking up the gain. Presumably the increase in gain would be general (reflected e.g. in muscle tone generally) but part of the "stereotyped response" is to give high priority to the fight/flight goal for the driver's seat. Obviously I'm groping around where I have no expertise, but the ideas in the article seem worth considering.

Perhaps the stereotyping sets up goals in conflict--fight OR flight--and thence the adrenalin and the high gain. There are other ways to increase gain more selectively, as meditators discover.

Semantics: I wasn't aware of supposing that words refer to things "out there" but I may have used conventional phraseology. I'll try to be watchful.

>>correlation of an outcome with a purpose or goal smells unscientific
>>from a conventional perspective

> such opinions are based on a misunderstanding of science

And that customary, conventional misunderstanding might be a stumbling block to communication, which was my topic.

variance vs. variation

OK, I wasn't in the audience for this, and shouldn't have responded. I have never wanted to learn statistics. (BTW, I am grateful to CT for confirming my intuition in this.)

>>if the work can be done within language . . . , why have recourse to
>>[an external] metalanguage

>I'm guessing that you're referring . . . to Avery's structural
>diagrams. . . . But can the work really be done within language?
>Doesn't it ALL have to be done within meaning? And isn't a structural
>diagram as good a way to refer to a meaning as a lineal sequence of
>words, and in some cases better?

If Avery's diagrams describe signals, connections, I/O functions, comparators, ECSs, constructions of ECSs, and the like, then they or diagrams using analogous conventions are to be used for all classes of perception, not just for language. That seems implausible.

But in fact they are diagrams to account for restrictions among words in their linear ordering within sentences. (See first quote from Avery, below--they're "just sentences.") They don't even get at discourse structures (other than a little bit of the more explicit cross-reference stuff that goes on between sentences). For that part of the work, operator grammar uses simpler means.

I'll have to return to the remainder of your (920318.0800) post at a later time, Bill. While I can steal a bit more time this morning let me segue if I can to . . .

Avery Andrews (920318:130615)

>A Jackendovian semantic structure is basically just a sentence with
>explicit & unambiguous head-argument relationships, plus some
>semantic typing.

The operator-argument relationships are explicit and unambiguous in operator grammar. The semantic typing is unnecessary, either because it is captured by words under reduction or because it is just wrong: the typing states or implies restrictions on word combination that it turns out people can violate given appropriate context. Of course I can't set out to prove that here, but much of what I have posted to the CSG list in the past shows how this works.

```
>The verb betray might have a lexical entry like this:
>
>  Category: V
>  Form:     /betre:/
>  Meaning: [event Does-sth-to(X,Y)
>                Y falls into the power of Y's enemies]
>
>(J has a fairly elaborate theory of verbal lexical entries, which I'm
>fudging here_.
>
>the noun `betrayal' would have the same meaning, but a different
>Form & Category, leading to differences in how the arguments get
>expressed.
```

The "Category" and "Form" fields are covered in obvious ways in operator grammar (I'll make that explicit if it's not obvious). The "Meaning" entry is very much like the "expansion" resulting from undoing reductions on sentences containing the word. But there would be no occasion to expand "betray" in this way in operator grammar. It would be a primitive element, an operator. Some usages betray (reveal) the reduction of other words, as for the metaphor in that last use and again in the following example (cited in the dictionary):

His best columns betray the philosophical bent of his mind.

(I've described the operator-grammar derivation of metaphor, so I won't make that explicit here, but tell me if I should anyway.)

"Event" and "someone does something to someone" I would take as coming from the nonverbal side rather than in the lexicon. To be sure, "someone does something to someone" is like the source of certain reductions, for example for product nominalizations, which goes roughly as follows (from memory):

Something which is a product of someone constructing the bridge
was imperfect ==>
The construction of the bridge was imperfect.

That is what differentiates betray and betrayal in operator grammar:

Something which is a product of John's betraying Fred,
was astonishing ==>
John's betrayal of Fred was astonishing.

For Jackendoff's
>   Y falls into the power of Y's enemies]
it is not at all clear to me why "fall into," "power," and "enemies" are more primitive elements than "betray." This sort of background knowledge for drawing inferences must come from knowledge associated with the word. As I said previously, there are several ways to represent this knowledge. Some that I know about:

1. In a set of discourses using the word "betray," including dictionary sentences stating things that are explicitly stated only to children or in fact in dictionaries and encyclopedias.

2. In the double-array representations of the information in those discourses, resulting from the application of operator grammar and sublanguage analysis.

3. In the union of those representations for each sublanguage domain.

4. In the nonverbal correlates of elements and relations in (3).

Such background knowledge is far too rich and variable (from one subject-matter domain to another, across speakers and across communities, through time) to be captured in a few formulae supposed to be the semantic decomposition of words. And as I said, the words in the formulae require further such decomposition on that perspective. Not so in operator grammar. All that work is to be done essentially on the nonverbal side, as reflected in (representable by) sets of discourses and presentations of the information-structures in discourses.

The example of dictionary sentences that I offered before was something like "and one uses umbrellas to stay dry." It may be that such sentences mostly or all state reference perceptions in explicit language. That leaves moot which is prior in acquisition or in precedence for control. Language in general is a way of representing nonverbal perceptions for our inspection and study, and possibly for our manipulation. I don't care whether these sentences expressing background knowledge are "basically" or "in origin" nonverbal or in language. By the correlation of nonverbal perceptions with word-perceptions, nonverbal perceptions become available for control as part of language, and things in language like assertions, injunctions, maxims, attributions, prohibitions, and instructions become available for setting nonverbal reference perceptions.

Gotta run. I've been at this over an hour and a half, and I've got a very full slate again today.

Bruce
bn@bbn.com

Date: Fri Mar 20, 1992 10:17 am PST
Subject: plummetting moths; drooling dogs

[From Rick Marken (920320)]

Reply to Avery Andrews (920320):

I said:

>>just get that damn signal outta here, sayeth the moth's
>>control system.

and Avery said:

>Well, that's exactly what I'm conjecturing that maybe is *not* happening,
>because the changes in the signal induced by the moths behavior aren't
>material to survivability of the moth. So my judgement so far would be
>that if what is going on is avoidance of being eaten via quick
>departure from the vicinity of the bat and the accompanying sonar
>signal, its control, otherwise it ain't. I'm not committed to its
>actually being one way or another - just to clarifying the issues.

By controlling the signal strength the moth generally does survive. From your perspective it looks like the moth is avoiding being eaten; that is certainly a way to describe the situation. But the moth probably knows nothing about that. Suppose that we set up a situation where we create artficial bat noises and, knowing that the moth controls them by dropping, we have another predator (instead of a nice beetle keeper) waiting at the bottom of the fall. Now the moth gets eaten as a result of controlling its sensory input. Are you saying that this means that the bat is not controlling in this situation -- because it is not avoiding being eaten?

I guess I'm not sure what you are trying to get at in the above paragraph. Are you trying to say that the dropping might still be a response to the stimulus sound even though the dropping changes the sensory effect of the sound? The only way this could possibly be true is if the sensor shuts down as soon as the moth starts to drop. Otherwise, the sensor output (p) is a continuous function of the bat's location (d) and the moth's position (o) so that

p = k.1d + k.2o (1)

since the falling (change in position, o) is clearly influenced by
the sensor output (p) we also have

o = k.3p   (2)

So the behavior os the moth is characterized by two, simultaneous
equations. Equation 2 is the SR law. I'm saying that you can ignore
equation 1 (which makes output depend on itself) only if you eliminate
the effect of p immediately after it is applied -- this would mean
bringing time into the equation. I suppose the equations would
then be

p = k.1d   t<s where s is the start of the fall and
o = k.3p   t>s

or something like that. But if you just leave things the way they
actually occur in nature then there is a closed loop -- no matter
what. Equations 1 and 2 apply simultaneously. We also know that behavior
is stable -- that means that the loop gain is <= 0. This means that the
coefficient in one of the equations must be 0 or negative. When we
get the signs of the coeffcients right and solve for o we get:

o = -k d

The output of a stable closed loop system depends on the disturbance
to the sensory input -- NOT on the sensory input itself. It looks like
SR (because d looks like a stimulus) but the moth does not respond to
sensory input (o is not a function of p).

I said:

>>You betcha!! Now think about how much you can "do" without influencing
>>one or another of your sensory inputs. When you realize that that amount
>>is almost precisely ZERO then you have entered the PCT zone.

> I see no reason why
>simple critters shouldn't have a fair number of these things rattling
>around in their circuitry.

I can think of no reason why they should. What could possibly be the
value of having inputs cause outputs that have no effect on the
inputs that caused them -- the outputs just going off into the world. Then
organisms would be truly like computers -- a classic SR device (usually).
You put stuff in at the terminal and that input gets churgled (I made it
up -- bet you know what it means) around and turned into output that
has absolutely no effect on the input -- unless there is a person there
to change inputs based on outputs.

Most organisms are built with sensors all over their bodies -- inside and
out. As I said, I can't think of any way that that body could be moved
without it having some sensory consequence. If you also imagine that
some of those sensory inputs are responsible for causing the body to
move in certain ways then you've got an organism that is locked in a
loop -- completely. It is possible that there are sensory causes of
outputs that are completely protected from any effects of those outputs.
Maybe there could be a connection between sensors on the tip of the

tongue and movement of the big toe. But why would there be many of such connections? Why do organisms generate outputs anyway -- just to give observers something to see? No, its to help them maintain THEMSELVES in some way. So if there are SR connections -- pure; no RS connection -- then the R must be occuring for a reason, right. Like to get the organism out of danger. I suppose that in a HIGHLY STABLE, DISTURBANCE FREE environment an organism could survive with some SR connections. The connection from S to R would also have to be HIGHLY reliable. So maybe you could have an organism that always makes response R when stimulus S occurs (and S is protected from the effects of R) and the effect of R on a variable that really matters to the organism can be counted on to be the same every time. But this has got to be a very rare arrangement.

I think that what happened is that evolution has created machinery (bodies) housing sensors that can be used to keep these sensor's outputs where they should be (as determined by the organism) in the context of the environments in which these bodies happen to find themselves. But what evolution (and "behavior" for that matter) is all about is sensory input -- PERCEPTION. Organisms exist only to keep their perceptions where they should be -- the outputs exist ONLY for the benefit of the senses. Unfortunately, from an observer's perspective, the process of controlling sensory inputs looks compellingly like like response to stimulation. Moving past this illusion is what PCT is all about.

These are GREAT questions Avery -- they really get to the nitty-gritty. Keep pushing and don't let us get away with anything. I think you are poking around where all psychologists should be poking around. We (PCTers) need psychologists to ask us these very tough questions -- because it get's right to the heart of what most psychologists think about PCT, and that is -- so what? We have to be able to answer the "so what?" question and we can't answer it in a convincing way unless people ask questions (or pose problems) like yours. Thanks.

Hasta Luego   Rick

Date: Fri Mar 20, 1992 11:54 am PST
Subject: Re: plummetting moths; drooling dogs

[Martin Taylor 920320 14:00]
(Rick Marken 920320)

Without wishing to disagree with Rick's conclusions that "Organisms exist only to keep their perceptions where they should be -- the outputs exist ONLY for the benefit of the senses", I do think that his argument is incomplete in respect of the dropping moth.

So far, we have: Moth senses bat sonar; moth changes mode of operation by folding wings and dropping; moth reduces sensation of bat sonar. But we don't have any sense of a feedback loop with measurable gain. We need the next stage: Moth's sensation of bat sonar in relation to a reference level leads to some change in error signal that results in some behaviour. Something here seems to me very S-R. No matter what the other circumstances (we hypothesise) an observer sees a sequence: bat sonar -> moth drop. There seems to be no disturbance on the bat sonar percept that would or could affect this (unless it be a bat-mute!). Even if we take it that the internal structure is an ECS that has as a reference signal zero-bat-sonar, the effect is still S-R-like. There would be more apparent control if the reference for bat sonar were non-zero, so that the moth sought out bats if it didn't get enough, but even in that case, a bat could CONTROL the moth if the reference level for bat sonar were fixed. It would still look like S-R for all practical purposes. And when the fixed reference level is zero (presumably plus some small increment to allow for low-signal detection problems) and the behaviour is always the same cessation of flying (or drop off a resting place?), there seems no way to discriminate between S-R and control.

The fact that the moth recommences flying when the sonar signal goes away can equally well be seen as S-R: Signal->response (stop flying and drop); no signal->no response. There is or isn't a feedback loop, depending on how you look at it. Certainly the moth's behaviour changes its environment, altering the stimuli to which an S-R person would say it responds. Certainly the moth's control system brings its percepts closer to their reference states. Who wins here?

If you have a predisposition to see all behaviour as control, you can easily describe the moth's behaviour as control. If you want to believe that some behaviour is S-R driven, you can see the moth's actions that way, and I do not believe that anything other than Occam's razor will say that you are wrong to do so. Will Occam's razor actually shave away the S-R interpretation? I don't know, but I have faith that it will.

Martin

Date: Fri Mar 20, 1992 12:10 pm PST
Subject: unusual occurrance; evolution and CT

[From Bill Powers (920320.0700)]

Avery Andrews and others in same boat:

To send copyright permissions for Closed Loop to Greg Williams, you can use e-mail. Greg's e-mail address is 4972767@mcimail.com.

David Goldstein (930218) --

Re: The Boy Who Pulled The Fire Alarm Because He Had a Feeling He Should And There Really Was A Fire.

When I lived in the Chicago area, I read that the Chicago Fire Department answered about 7000 false alarms per year (yes, 20 a day). This astonished me, because in my entire life I have never seen a fire engine called out on a single false alarm (that I knew about). Actually the number could be even higher than the Fire Department recorded, because with so many false alarms in Chicago, or New York, or Detroit, or any other major city, the firefighters must occasionally have actually found some sort of fire in the vicinity when they got to the location of the false alarm, and didn't realize it was false. I wonder in how many of those cases the prankster was quick-witted enough to step forward and claim credit for having known about the fire, or to claim to have had a feeling about it (even more interesting).

The problem with being witness to rare events is that a single person gets an entirely wrong idea of how frequent such events really are. As far as my personal experience goes, knowing someone who had a forewarning of winning big in the lottery and then won is so rare that there's no explaining such a phenomenon. But there are many other people who know someone who just had a feeling they had better buy a lottery ticket, and they won! Every week there is a big winner in many states, and among these winners you quite often hear stories about premonitions, hunches, and so on that led them to buy that one extra ticket which was a winner. "How do you explain that?" they demand. I explain it very simply: the lotteries are designd so somebody is almost sure to win the big prize, and when many people believe in luck, hunches, magic numbers, premonitions, precognition, guardian angels, favors from saints, answers to prayers, effectiveness of talismans, and holding the thought, the chance of a winner having had some sense of forewarning is actually quite high. Of course hundreds of millions of others who had similar forewarnings, week by week for years, did NOT win. I have known many of them. But who remembers them? They don't even remember losing themselves; they go out and buy another ticket, sure that they have a terrific chance now that they've lost exactly 199 times (base 10) in a row. The will to believe is indomitable.

That young man who got caught pulling the fire alarm and played innocent was pretty lucky -- there really was a fire. I wonder how many hundreds of false alarms this boy was responsible for up to then. I'm sure he didn't tell anyone about them, especially his parents. Far more interesting to say "Golly gee, I don't know what came over me, something just told me I had to pull that alarm." Clever little psychopath, but then, weren't we all?.

I have a hunch, however, or something tells me, that this explanation of the "unusual occurrance" is not going to satisfy everyone. I expect to get a modest trickle of character references for the boy, especially from people who never met him. You see? Even I have precognitions now and then. Eldritch.

The Plummeting Moth, the Drooling Dog, and models of evolution:

What distinguishes S-R theory from control theory in matters like this is that S-R theory claims no special relationship between the stimulus and the response, while control theory claims that the response is really an action aimed at controlling the perceived stimulus. The bat emits the sonar chirp; the sound energy reaches the moth, the moth obediently folds its little wings and drops like a stone, and according to S-R theory that's the end of it. Of course, it is explained, there is a reason for this: that response is selected for because moths that exhibit it survive somewhat better than moths that don't, and because moths compete within their niche, the better survivors come to predominate. The logic is quite airtight even if there are some missing details.

If the S-R/evolutionary explanation is correct, then we should be somewhat surprised to find that a response affects a proximal stimulus before it has finished occurring (even if it doesn't affect the remote cause of the stimulus). We should be puzzled to find that even the first moment of proximal stimulation is affected by whatever behavior is going on at the time. If we believe in cause and effect, it would be worrisome to think that a regular remote stimulus has a proximal effect on the sensors that is anything but regular, and in fact depends very strongly on what the organism is doing, how it is oriented, how fast it is going, how much effort it is already exerting, its internal state, and what else is independently affecting those same proximal stimuli at the same time. Why is there a regular dependence of the behavior on the remote stimulus, when the sensory effect of the remote stimulus is so strongly affected by behavior, even at times to the point of reversal?

Control theory makes a prediction that S-R theory doesn't. It predicts that whenever a regular response (consequence of motor actions) occurs, that response affects the same sensory inputs on which the apparent stimulus acts. Of course this isn't quite as clear-cut a prediction as it seems, because some research may be needed to find out what those sensory inputs are that are affected both by the remote stimulus and by the response, and that end up being stabilized by the combination of remote stimulus and response. This slight difficulty is compounded by the fact that the apparent remote stimulus noticed under S-R theory may be only loosely or indirectly related to the proximal stimulus that matters -- the remote stimulus is defined by the interest the human observer has in it, not by its importance to or even perceivability by the organism under study. So to check out this prediction from control theory, it may well be necessary to return to the original phenomenon, looking for something that the S-R theorist wouldn't ever have noticed. That's the trouble with trained observers: they're trained to see what their theories lead them to expect, and to ignore irrelevant details. What's irrelevant under one theory may not be irrelevant under another.

This problem has faced not only evolutionary theory but reinforcement theory as well. Nothing in reinforcement theory says that the action being reinforced has to have any particular effect on the current reinforcement process. There's no reason why a hungry animal should, for example, eat instead of drink or run in a wheel when reinforced with food. To say that the animal eats because it is hungry is to put a cause inside the animal, and this goes contrary to the basic concept of behavior being directed by external events (to which many biologists subscribe). The fact that food reinforcements increase eating behavior and not some other kind of behavior can be explained only on evolutionary grounds; organisms that react this way survive, while those who don't have fallen under Darwin's Hammer.

Evolution, presumably, is in some way responsible for the fact, under control theoy, that organisms control their own inputs. If this is the case, what is it, roughly, that evolves in a control-system type of organism? Not the actions, because the actions vary with every passing disturbance of the inputs, cancelling the effects of the disturbance. Not the relationship of the disturbance to the actions, because that is determined completely by the fact that a closed-loop organization is present; if a control system exists, there is no possibility of variation in that relationship. All that is left is the physical organization itself consisting of a perceptual function, a comparator, an output function, and an inner reference signal. If, as seems plausible, the inner reference signal is derived from past values of the perceptual signal, evolution can't affect the reference signal directly: it is either some value that has been perceived before, or zero, or random.

So we are left with the inheritance of perceptual functions, comparators, effectors, and memory systems. These functions may be simple or complex;

effectors, for example, may produce temporally-patterned outputs in response to their inputs. What cannot be inherited is the particular action that will be produced, how the effector will be driven: that must remain dependent on events in the current environment and the settings of the internal reference signals that arise from experience.

There may be, in lower organisms, preformed functions that are not later modified through interaction with the environment. This could even include preformed memory systems containing reference signals implying specific experiences. The bower bird, apparently, inherits a sketchy reference image of a physical structure, the bower. It's up to the present organization to find a way to match that reference image with a present-time perceptual image, so no two bowers are alike in the way the structure is created. The movements and detailed control processes that bring the bower into being are clearly not inherited because they take advantage of available materials, such as fragments of brake-light lenses, which evolution surely could not aniticipate being present.

So under control theory, it is conceivable that perceptual functions and reference images may be inherited, and that some output function organizations may be inherited. What is not conceivable is that behavioral outputs can be inherited, because if those behavioral outputs are fixed, they will not be able to vary as required by the circumstances of the current environment, which are unpredictable.

My impression of evolutionary models is that they begin with the basic premise that behaviors are inheritable, or at least S-R connections are inheritable, which is almost the same thing. Given that the assumption is right, that genes specify responses, the rest follows. By making survival contingent on performance of certain responses under certain external conditions, and given some mechanism for randomly reshuffling or mutating the genes which are by fiat attached to response measures, one can construct a model in which the surviving individuals will come to exhibit the required responses. The kind of reshuffling processes that is used will determine how rapidly the unfit are weeded out, but the ultimate result is inevitable. The nature of the final behavior is determined by the rule connecting behavior to survival.

If, however, responses to stimuli are not in fact inherited, this kind of evolutionary model loses its relevance: it does not describe a real organism but only a possible one. The genes cannot specify specific responses to specific stimuli. If the organization that best serves survival is a control system, then an evolutionary model must be cast in terms of inheritance of control-system properties, not behaviors.

Organisms with similar control-system properties may behave similarly in similar circumstances, but that is not because they inherited similar behaviors. It is the result of a control-system type of organization interacting with a physical environment that creates particular relationships between effector outputs and consequent effects on sensory inputs. People open doors by turning the knob and pushing or pulling not because they have inherited those moves, but because doors are constructed to open that way. If a door has stiff hinges, the person will pull harder, just enough harder to make the door open in the usual way. This is not an inherited response to stiff hinges. It is the basic mode of operation of any control system: a bigger error produces a bigger output.

When we introduce learning or reorganization, even the structure of inherited control systems loses its connection with inheritance. The behaving system now has the ability to alter its own structural organization, provided only that evolution gave it enough preorganization to start with, including the capacity to reorganize. The survival-selection criteria now must have to do with the adequacy of the beginning organizatiion and the efficiency of the autonomous reorganizing system. As this is probably the condition in which we find human beings, the idea of a adult human being's behavior being inherited no longer seems even remotely feasible.

I'm sending a copy of this to Randy Beer as well as to CSGnet.

Best to all,

Bill P.

Date: Fri Mar 20, 1992 12:11 pm PST
Subject: Arm model; falling moths; HCT language model

[From Bill Powers (920320.1100)]

Wayne Hershberger (920319) --

Here's the diagram I use for a pair of muscles (somewhat improved). I lay them out in a straight line although they really lie along opposite sides of a bone and work across opposite sides of a joint, like a pulley.

```
                            alpha signals
MUSCLE MODEL            a1                  a2
                         |                   |
                         |                   |
                    ---------           ---------
Anchor--////////---|    ---------Load--------    |---/////////--Anchor
                    ---------           ---------
        Elastic    Contractile         Contractile    Elastic
```

I consider the differential signal, a2 - a1, to be the composite driving signal. When a2 increases and a1 decreases, the movable central parts of the contractile elements both move to the right. The springs do not change length. Thus the load moves to the right. There is no applied force after the move.

With a2 > a1 the configuration looks like this:

```
MUSCLE MODEL            a1 - delta          a2 + delta
                         |                  |
                         |                   |
                    ---------           ---------
Anchor--////////---|       ---------Load-------- |---/////////--Anchor
                    ---------           ---------
        Elastic    Contractile         Contractile    Elastic
```

A force applied to the load (the "load point") to the right relaxes the right spring and stretches the left one, creating a restoring force to the left toward the undisturbed position of the load. This leftward force is the muscle force opposing the force applied to the load. I assume that the force does not alter the configuration of the contractile part; only the signals determine how far in the plungers are.

The control systems using this pair of muscles are shown this way; the spring is zero-centered:

```
                       Alpha
                     Reference                                 Gamma
                         |             length gain           Reference
                --<- comparator <----------Kg-------             |
               |        ^       <-----Kd----   length error      |
          error|        |           damping |    sensor          |
           sig |        |       d(length)/dt|   ----------       |
               |        |                   |  |          |      |
       output  |     force gain             |<-|Mechanical|<-----
         gain  Ko       Kt                     |Comparator|
               |        |   Force sig          |          |
               |        -------                 ----------
               |              |                      ^
               V    Spring    |tendon                |
            ------      Ks    |sensor                |
    load ------   |--///////---O-----Anchor          |
            ------                                   |
     |<-------------------------------->|            |
                    |                                |
                     ----------muscle length--->-----
```

The effective spring constant with high loop gain, is Ks/KoKt - Kg/Kt. It can thus be zero, if Kg = Ks/Ko. The apparent mass is Kt*Mo, where Mo is the actual mass (or moment of inertia). I misspoke myself when I said the apparent mass was reduced. Typical values used in the model are

Kg = 50 NSU/radian
Kt = 50 NSU/newton
Ks = 50 newton/radian
Ko = 1 newton/NSU
Kd = 0 to 10 NSU per radian/sec

"NSU" means "nervous system unit." "Lengths" in the above diagram are converted to radians of movement about the joint.

To answer your question, the values of the gamma inputs for each control system are set to 0 units in azimuth, 0 in elevation, and 512 in elbow flexion (the units are such that 4096 angle units = 360 degrees).

When gravity is turned on, with the alpha reference inputs set to zero (a mode of the model available for testing), the arm slowly sags as if sinking through heavy molasses. The actual positional loop gain is quite low. I don't think that the stretch reflex loop is actually a position control system; its evolutionary purpose seems to be to alter the apparent arm dynamics, mass, and spring constants to make control by larger loops (possibly involving the joint angle receptors) very stable and independent of limb segment interactions. I don't really understand why it works as well as it does. The bare mass-spring properties of muscle and arm are drastically altered by the control system parameters when set for the best performance.

It will be interesting to see what happens when I put in the sixth-power muscle spring nonlinearity, the force-velocity dependence of the muscle, and the changing mechanical advantages as the joint angle changes. Version 3. I should also put in the different response of the muscle to onset and offset of signals, and the branches of the biceps and triceps that span both shoulder and elbow joints. All this really requires modeling the opposing muscles separately. Maybe Joe Lubin and his students, plus Greg Williams, would like to carry this on to version 3. I want to take a vacation from this arm model now and get some other things done (after writing a paper with Greg for publication).

Avery Andrews (920319) --

>So my judgement so far would be that if it what is going on is >avoidance of being eaten via quick departure from the vicinity of the >bat and the accompanying sonar signal, its control, otherwise it ain't.

My judgment would be just the opposite, because I don't think a moth can perceive "being eaten" or "departing from the vicinity." The moth has to get by on perceiving and controlling some fuzzy blobs of light, some smells, and some sounds of variable intensity and possibly direction relative to its body. It can't possibly make use of the information that controlling these perceptions will help it survive, or even that it's being threatened by a "bat" and all that this implies to a human being. It would need a brain the size of yours to control for such things.

The moth doesn't behave as it does because that behavior is material to survival. Cause and effect run the other way. The moth survives because it controls for the variables it can perceive, with respect to the reference levels it uses. Controlling a certain sound relative to a low or zero reference level is apparently enough to permit its survival (usually, or sometimes) when in the vicinity of bats; however, it has no idea what that sound means, or that the consequence of this behavior is "survival." It's just a bad sound, to be avoided. Neither does it have any idea that it's falling through space when it closes its wings. That's a human perception. It just makes its wings feel a certain way and that suffices to control the level of the sound, to a sufficient degree. The moth can't know anything about the details of why this works. Nor does it need to, to manage its little world as well as a moth can. We could explain to the moth why controlling for just those variables in just that way is a very good idea for the moth, but the moth wouldn't understand.

I imagine that we would look much the same to a vast cool intelligence from an advanced civilization elsewhere. We control what we perceive, and we survive, but we don't really know why this works.

Bruce Nevin (920320) --

>Language in general is a way of representing nonverbal perceptions for >our inspection and study, and possibly for our manipulation. I don't >care whether these sentences expressing background knowledge are >"basically" or "in origin" nonverbal or in language. By the >correlation

of nonverbal perceptions with word-perceptions, nonverbal >perceptions
become available for control as part of language, and >things in language
like assertions, injunctions, maxims, attributions, >prohibitions, and
instructions become available for setting nonverbal >reference perceptions.

Frame this and hang it on the wall. I think this statement consolidates a
large amount of progress toward a coherent HCT theory of language. If you
keep talking like this, you can expand as many reductions as you like and
I'll remain meekly silent. You're starting to define the problem. The
better you can define it at this level of discourse, the quicker all those
messy details will fall into place.

Here's a thought to chew on. Part of the problem most people have in
communicating and/or thinking is their very failure to see that the
meanings communicated by the words they produce and hear are incomplete and
ambiguous. If they tried to expand all communications, either your way or
my way, they would discover how bad is the information they're receiving
and emitting. "I'm glad you agree with me," they would say, "but what,
exactly, are you agreeing with?"

Best to all

Bill P.

Date: Fri Mar 20, 1992 12:14 pm PST
Subject: Re: Misc responses

[Martin Taylor 920320 14:30]
(Rick Marken 920319)
(I don't know why I am doing this, with so many deadlines so close...but
this is more fun. Conflict...resolution)
>
>Martin Taylor (recently) says:
>
>> If there are two schools of thought, each with
>>good reason claiming that they have the truth and the other doesn't, they
>>are probably both right, except in that claim.
>
>Well, I never claimed to have the truth. I claim that HPCT is the best current
>explanation of a phenomenon that is currently not studied in psychology --
>except obliquely -- control. So I claim that that model is better than
>any other as an explanation of that phenomenon. I guess I am also saying
>that most of what psychologists think of as the phenomenon called "behavior"
>is actually control -- so I am saying that my model is the correct model
>of behavior (as far as we can go towards correctness with current know-
>ledge) and their models are, thus, wrong.

Yep, that's the attitude I was taught about in grad school. Think about
how much can be properly treated by "classical" methods even if their
model is "wrong." Think about the fact (about which we are believed to
agree) that a control system with a fixed reference works like an S-R
system. Think about the many reference levels that can be assumed to be fixed
by the agreement of the subject to participate in an experiment. And then
think carefully about whether their "wrong" models might not be a subset
of your model, even if neither they nor you find it intuitively to be so.

I'm not claiming you are wrong, nor that "they" are. But I have learned to be wary about the fragility of structures that fit together in lots of ways. Even if the foundations are insecure, the superstructure often holds together when they are replaced by better foundations. Newtonian mechanics does not fail because we now understand that kinematics is all a question of geometry, not mystical "forces."

>>It is obvious to me that the basic ideas of PCT have to be right, just as
>>are those of (say) signal detection theory or information theory.
>
>I think I don't agree -- maybe I don't understand.

All I'm saying is that you can't legislate pi to be 3. To disagree that the basic ideas of signal detection theory or information theory are correct is equivalent. It is only an opinion, though, that the basic ideas of PCT are right. I hold that opinion, but it could be changed by experimental evidence. Pi cannot be set to 3 by experimental evidence.

Martin

Date: Fri Mar 20, 1992 12:40 pm PST
Subject: Re: Taylor catchup

[Martin Taylor 920320 14:45]
(Bill Powers 920319 09:00)
>
>>From what I remember of your previous remarks about zero references, the
>problem seems to come up because of thinking in digital rather than analog
>terms. In digital terms, a high-level variable is there or not there, so
>the error is either present or absent. If you look at any particular
>example of such cases, you can see how the apparently digital variable can
>be subject to analog disturbances. Elements of the perception aren't just
>right or wrong: they can be almost right, or a little wrong. These
>differences call for variations in lower-level reference signals to keep
>them from getting large enough to constitute a serious error.
>
No. Digital vs. analogue has nothing whatever to do with my argument. The argument has to do with analogue ECSs that are nearly linear in the vicinity of their control reference point. Dead zones and one-sided controls change the argument, and it wouldn't apply at all to digital systems (I think).

The hypothesised situation is one in which distrubances have been small and slow enough in the past to allow some part of a control hierarchy to stabilize. If all the ECSs in a level are orthogonal, then all their percepts must be matching their references, and hence they are emitting zero error signals, which form the references for lower levels.

Naturally, the environment will be disturbing this idyllic peace, but by hypothesis the disturbance is slow enough that control can be maintained very closely.

Rick pointed out that if the ECSs are not orthogonal, then there is a residual tension that means that the error signals are not zero when things have stabilized. This is correct, and was a point I wanted to make in further discussion. All the same, this non-orthogonality effect applies within a level, so far as I can see, and does not prevent lower (orthogonal) levels from attaining zero error that is provided as reference for yet lower levels.

You pointed out that if the sequence level is involved, the reference signals sent to lower levels are always changing, and therefore the low parts of the hierarchy cannot stabilize. This changes the hypothesized conditions. But in itself it led to the question we now recognize to be unsolved, as to what is the error signal at the sequence level, and when is it expressed.

As far as I am concerned, that's where the discussion lies, and is the starting point for what I would like, someday soon I hope, to expand on: Situation awareness and workload assessment.

Martin

Date: Fri Mar 20, 1992 2:32 pm PST
Subject: Steam Engine - RKC

March 20, 1992

Subject: Steam Engine - RKC

[From Robert K. Clark (491-2499)]

Chris Malcolm and the Watt Steam Engine Governor.
Sorry to be slow to respond. My time for SCG is very limited as I am responsible for several on-going activities. My CSG files are still not well organized. There are many interesting topics, but I can't follow them all!

I'm not sure what you are suggesting with this. Are you suggesting that it is not "possible to trace all the signals/events through the system??" Or that it may be difficult to do in some cases? This can be particularly hard to do if digital computers are involved in the absence of the system documents. With analog systems, or other hardware systems, it may be easier.

What I had in mind was the tracing of each step in the system in cause & effect terms. This is what sometimes results in tracing the signals ("events") around and around the loop, with the reference signal ("set-point") fixed. And then tracing the systems, likewise, for the effects of changing the set-point both with and without having the out-put "signal" fixed. Tracing the events through the hardware may be awkward -- this is one advantage of the Control System terminology. It has broad general applicability and is convenient in many situations.

Of course much of the interest in Control System Theory of Behavior is concerned with application to living beings, mainly, people. And this quickly moves to tracing neural networks. Here it is not uncommon to find people going "around and around the loop."

Perhaps my following response is unnecessary, but here it is:

The Watt Steam Engine Governor is indeed a negative feedback system of brass and steel, with an adjustable set-point. Note I said "may be possible to trace all signals/events through the system." These "signals/events" need not involve conversion to alternative physical form. After all, the concept of "signals" is just that, a "concept."

In the Steam Engine situation, there are several ways to describe the "signals." For example, the distance of the spinning balls from the axis can be regarded as the "feedback signal." This distance is then converted by levers into a force applied to a movable valve, countered by an externally adjustable tension on a spring or other opposing force. As the valve opens (or closes) the engine operates to change the speed of the assembly, particularly the rotating ball system. Thus each step in the operation of the system is of a "directly cause and effect" or "stimulus response" nature, while the over-all operation serves as a control system. Note, however, that, without some entity to adjust the "set-point," the operation of the Steam Engine as a control system is incomplete.

There are many non-living systems in which the physical expression of one or another signal may be different from what is expected. For example, many thermostatic systems detect temperature by means of the shape of a bi-metallic strip. In such a case, the temperature signal can be regarded as the position of the end of the strip. Which may, in turn, make or break an electrical contact -- the comparator -- producing an output signal to some part of the output function.

The physical expression of some feedback systems may be difficult to untangle, but if the conditions defining such systems are met, the over-all operation is clear.

You speak of "forecasting." Indeed, you are right. In some situations a great deal can be accomplished without detailed information. Meteorology, for instance. There are many examples of this in physics and elsewhere. What I don't like about Determinism (hence "Behaviorism" carried to the extreme) is the implication that one is powerless to change ANYTHING! Of course this is extreme, and usually handled by limiting the application of the deterministic philosophy. I say I "don't like Determinism" because that is essentially my orientation. I have been unable to find any way to demonstrate either determinism (complete, unlimited) or free will (limited, of course). So I "takes my choice."

I reacted to the quotes from Bill's remarks about BEHAVIORISM, with which I agree, because it seems to me that a more philosophical viewpoint is more effective. This also begins to show my approach, primarily that of an experimentalist who needs a working theory -- that works -- in order to interpret and communicate his experimental results.

Regards, Bob Clark

Date: Fri Mar 20, 1992 3:10 pm PST
Subject: Re: plummetting moths; drooling dogs

[From Rick Marken (920320b)]

Martin Taylor (920320 14:00) says:

>Without wishing to disagree with Rick's conclusions that

Not to worry. I'm used to it. Besides, I just took a brief foray this week into one of the network newgroups. The tone was so different -- angry, hostile, insulting -- compared to CSGNet it was amazing. A real eye opener. I learned nothing except never to do that again. And the amount of real information that gets transmitted here relative to there is astounding. It's incredible what you can accomplish when the goal is cooperation rather than victory. Gary -- I second my vote to keep us from becoming a newgroup. Even at the height of the "Beer bash" we looked like a church social compared to what goes on in those newsgroups.

>So far, we have: Moth senses bat sonar; moth changes mode of operation by
>folding wings and dropping; moth reduces sensation of bat sonar. But we
>don't have any sense of a feedback loop with measurable gain. We need the
>next stage: Moth's sensation of bat sonar in relation to a reference level
>leads to some change in error signal that results in some behaviour.

I think not. The equations that I gave in the Avery article have no explicit reference input -- there is no theory of the organism other than an SR theory -- o = kp -- but when you solve them simultaneously you end up with the basic facts of negative feedback control:

o = kd and

p = 0

(because there is no reference, with high gain the sensory input is driven to the 0 of the perceptual variable scale.)

There are some assumptions about gain that I made -- but they can be based on inspection. The gain is not 0 because there is output to input. The gain is negative because, if you get the signs of the coeficients so that gain is positive then the equations make no sense as descriptions of behaivor.

Bill P. did all this stuff much more elegantly and beautifully in the Psych Review article (1978). He basically showed that, with negative feedback from output to input there IS control -- and an SR model of the situation is simply WRONG.

>Something here seems to me very S-R.

Yes, equation 2 describes the organism as an SR system. The FACT of control is derived from recognition that there is also an RS law and that taking both facts into account simultaneously reveals that the system is CONTROLLING.

> No matter what the other circumstances (we
>hypothesise) an observer sees a sequence: bat sonar -> moth drop. There
>seems to be no disturbance on the bat sonar percept that would or could
>affect this (unless it be a bat-mute!).

The disturbance is to the sensory variable that causes the moth drop. One major disturbance to this variable is movements of the bat -- think of my d as the effective intensity of the sound, which depends on how far the bat is from the moth at any instant. As the bat moves, d changes and, hence, p changes.

>   there seems no way to discriminate between S-R and
>control.

There is -- but it's not simple (obviously -- or everyone would have noticed by now).

>If you have a predisposition to see all behaviour as control, you can easily
>describe the moth's behaviour as control. If you want to believe that some
>behaviour is S-R driven, you can see the moth's actions that way, and I do
>not believe that anything other than Occam's razor will say that you are
>wrong to do so.

Control can be tested -- and proved to exist. The difference between SR and control is BIG and needs no razor for choice. It just requires studying the equations, doing the tests and looking at what is going on. I think the continued predilection for seeing behavior in SR terms is purely psychological -- it's easier and more familiar and it takes little (or no) learning to understand it. I think eventually control will be accepted -- just as the fact that the earth spins on its axis is now accepted. All the obvious evidence is against both. I think people eventually accepted the spining earth just because all the experts said so (of course, when we could take pictures of it that made it even easier). It didn't take occam's razor to make Copernicus right -- it took tough, detailed thinking. Most people will not want to do this -- so when control theory is finally accepted in the life sciences, people will just accept is unquestioningly (as they accept a spining earth) and wonder how people could have been so silly as to believe in SR -- they will even forget that the SR view (like the flat stationary earth view) is the one that is by far the more obvious.

Hasta Luego   Rick

Date: Fri Mar 20, 1992 4:26 pm PST
Subject: Re: plummetting moths; drooling dogs

[Martin Taylor 920320 18:40]
(Rick Marken 920320b)
I didn't follow your post to Avery on the moth feedback loop, and I don't follow your rebuttal to my claim that one could see the moth's behaviour in S-R terms. Here's the quote and problem.
>
> Are you trying to say that the dropping might still be
>a response to the stimulus sound even though the dropping changes
>the sensory effect of the sound? The only way this could possibly be
>true is if the sensor shuts down as soon as the moth starts to drop.
>Otherwise, the sensor output (p) is a continuous function of the bat's
>location (d) and the moth's position (o) so that
>
>p = k.1d + k.2o (1)
>
>since the falling (change in position, o) is clearly influenced by
>the sensor output (p) we also have
>
>o = k.3p   (2)
>
>So the behavior os the moth is characterized by two, simultaneous

>equations. Equation 2 is the SR law. I'm saying that you can ignore
>equation 1 (which makes output depend on itself) only if you eliminate
>the effect of p immediately after it is applied -- this would mean
>bringing time into the equation. I suppose the equations would
>then be
>
>p = k.1d   t<s where s is the start of the fall and
>o = k.3p   t>s
>
>or something like that. But if you just leave things the way they
>actually occur in nature then there is a closed loop -- no matter
>what. Equations 1 and 2 apply simultaneously. We also know that behavior
>is stable -- that means that the loop gain is <= 0. This means that the
>coefficient in one of the equations must be 0 or negative. When we
>get the signs of the coeffcients right and solve for o we get:
>
>o = -k d
>
>The output of a stable closed loop system depends on the disturbance
>to the sensory input -- NOT on the sensory input itself. It looks like
>SR (because d looks like a stimulus) but the moth does not respond to
>sensory input (o is not a function of p).
>

In (1) you have the sensor output as being the sum of two linear functions of the respective locations, whereas I think it should be a non-linear function of the vectorial difference of the two positions with terms in the bat's orientation relative to the moth. The appropriate form should probably be something like (ignoring the bat orientation)

p = f(o-d) (1a)

and (2) should be in terms of the derivative of o, not of o itself (actually the derivative of the z component of o would be closer, but let's not worry about that). So I think we should have

do/dt = g(p) (2a)

The moth "wants" p to be minimized. These equations in themselves give no hint as to how that can be done, so far as I can see, since (2a) is really, as you say, a statement of the S-R behaviour, and whether dropping will increase or decrease p depends on whether the moth is above or below the bat's beam.

The equations cannot be solved, because integrating (2a) provides a constant of integration that depends on the initial value of g(p). g(p) itself is a problem, because it probably is discontinuous, switching between 0 and some fixed value at some small value of p. All we can get, so far as I can see, is

do/dt = g(f(o-d))   (3a)

which is a statement of the S-R "event" in externally observable variables.

I don't see any demonstration of control here, still. I see consistency with control, and consistency with an S-R view. The only reason, still, that I see for preferring the control view is that it is more parsimonious, and covers a wide variety of situations including this, whereas the S-R view can be applied only under special circumstances, all (?) of which can be seen as degenerate instances of control.

As Bill has often said, if the world is stable, it is possible for the same actions to have the same effects most of the time, and you don't NEED control, even if the opportunity for control is inherent in the situation. Whether control is actually exercised when the opportunity is there is a matter for experiment.

I suppose that moths that drop into the bat's path rather than out of it get eaten, but on average moths that drop get eaten less often than moths that ignore the bat. They don't know it, as Bill says, but Nature knows whether they have offspring.

Martin

Date: Fri Mar 20, 1992 5:06 pm PST
Subject: S-R like control theory???

[From Bill Powers (920320.1800)]

Martin Taylor (920320) --

Help! You're running away with the ball.

>Think about the fact (about which we are believed to agree) that a
>control system with a fixed reference works like an S-R system.

Consider:

```
Distal stim --f1---> proximal stim -------> Organism ---> action
                           |                                 |
                            ----------neg feedback ---------
                                      function, f2
```

The above is not a matter of theory or choice. If the negative feedback path to the proximal stimulus exists and the gain (action/proximal stim) is high, then the action will be very nearly determined by f2 and f1, rather than by the organism's properties. The existence of the negative feedback function and the remote-proximal function can be established by inspection of the environment -- no theory involved. If the above relationships are found, THIS IS NOT AN S-R SYSTEM AND IT DOES NOT WORK LIKE ONE. This is quite aside from the setting of the reference signal.

"Proximal stimulus" is just a word for the observable counterpart of a controlled variable, when control is present.

An observer's assuming that there is an S-R relationship between distal stimulus and action will result in attributing the form of the observed functional relationship to the organism: the observed relationship is action = inverse(f2(f1(distal stim)), or f2/f1(d.s.) in the linear case. In fact, the form of the observed relationship is fixed by the environmental functions and in no way characterizes the organism. So the S-R theorist will think that behavioral laws are being discovered when they actually have nothing to do with the organization of behavior.

S-R theory is a subset of control theory only in the sense that the control model explains why particular stimulus-response laws are observed (when control is in fact present). At the same time, control theory shows that these laws are illusory; they describe not the organism but its environment.

So there is very little in the S-R behavioral literature that actually increases our knowledge about the properties of organisms; S-R formulations are not just somewhat simpler or alternative descriptions of the same thing that control theory describes. They are, when control is actually present, misrepresentations of the organization of behavior.

Best   Bill P.

Date: Fri Mar 20, 1992 9:12 pm PST
Subject: Nobody does it better

[Rick Marken (920320c)]

I've really GOT to get a life (as my daughter says).
Anyway, here I go again. Breath easy; it will be short.

Martin Taylor (920320 18:40) says:

>I didn't follow your post to Avery on the moth feedback loop, and I don't
>follow your rebuttal to my claim that one could see the moth's behaviour
>in S-R terms.

[Lots of complex math deleted]

>I don't see any demonstration of control here, still. I see consistency
>with control, and consistency with an S-R view.

Just as I was about to ask my son to teach me calculus (he's the
math genius -- just my luck, he has zip interest in control theory)
I receive this:

>[From Bill Powers (920320.1800)]

>Martin Taylor (920320) --

>Help! You're running away with the ball.

```
>Distal stim --f1---> proximal stim -------> Organism ---> action
>                            |                                |
>                             ----------neg feedback ---------
>                                        function, f2
>
```
>The above is not a matter of theory or choice. If the negative feedback
>path to the proximal stimulus exists and the gain (action/proximal stim) is
>high, then the action will be very nearly determined by f2 and f1, rather
>than by the organism's properties. The existence of the negative feedback
>function and the remote-proximal function can be established by inspection
>of the environment -- no theory involved. If the above relationships are
>found, THIS IS NOT AN S-R SYSTEM AND IT DOES NOT WORK LIKE ONE. This is

>quite aside from the setting of the reference signal.

Now why didn't I do it that way. Imagine the savings in bandwidth.
When I saw this, despite the fact that it is Bach's birthday
tommorrow, the crass but haunting strains of the pop song
"Nobody does it better" started to waft through my head.

Thanks Bill. Regards Rick

Date: Sat Mar 21, 1992 12:17 am PST
Subject: plummeting moths

Regrettably, I don't have time to respond carefully to everything that has been written about plummetting moths, since I have to get ready to go off to Palo Alto in a week, but here's how I see things so far.

S-R psychology is a dead-end because it misrepresents what's going on in at least three interrelated ways:

a) the nervous system is seen as a discrete transducer over an utterly ill-defined class of event-types (stimuli and responses) rather than a continuous transducer over vectors (I'm assuming that the fact that its actually pulse-trains is irrelevant for the time-scales relevant to behavior).

b) the functional importance of the R -> S link is ignored

c) it isn't noticed that much of what is going on is the stabilization of the values of complex functions of inputs rather than the production of transient effects under given circumstances.

And, just as Gary, Bill et. al. say, even though S->R psychology is officially dead, it is still seriously alive as an influence in AI, philosophy of mind, etc. (I guess I differ from Gary about Phil of Mind in that I think there's plenty stuff to be salvaged from there).

(a) is an erroneous description of the psychological mechanism (hopelessly vague except for the bit about discreteness, which is false), while (b) and (c) are basically mistakes in ecology. By which I mean: suppose someone were to ask why this neural circuit is hooked this way rather than some other way? You can't (in most cases) give a sensible answer without including the R -> S relation in the story, and explain about how negative feedback loops work.

Now, considering my version of the moth (the one that crashes into the leaf litter), I continue to insist that this is an S-R setup rather than control, because, w.r.t. (a) in this particular case, there really is a discrete stimulus producing a discrete response [but N.B.: the stimulus is not a discrete *event* but a discrete *condition*, e.g. something that becomes true and stays that way for a while, and similarly for the response--I'm guessing that this possibility is catered for in the standard S-R conception]. (b) even though there is an R->S connection, it does *not*, in this particular case, contribute anything to the selective advantage conferred by this piece of gadgetry (except in the peculiar negative sense that if the action that caused plummetting also caused the sonar signal to become inaudible, without masking the moth from the bat, the setup wouldn't to its job). (c) this particular system does not in fact stabilize any perceptual function in the moth.

None of this means that S-R analysis is a good way to look at behavior in general, in fact the reverse, as can be seen by considering all the very specific conditions that have to be satisfied for this particular arrangement to be useful. It works because the moths happen to be in an environment whereby a very simple maneuver will take make them invisible to bat sonar, and not too vulnerable to other dangers (it will fail and thus not evolve in environments where the ground is normally swarming with ants hungry for hunkering moths, and equally fail in swamps).

It seems to me that there can be no harm for CT in recognizing this kind of thing as an S-R setup, and going on for a bit about all the rather special circumstances that have to exist in order for it to work: it's suitable only for some uses by critters who can't afford much in the way of brainpower, and rely on mass reproduction rather than individual survivability. Marginality is not the same thing as nonexistence, and looking at the properties of real cases where S-R setups can work ought to make it even clearer why, most of the time, they don't.

Avery.Andrews@anu.edu.aucl

Date: Sat Mar 21, 1992 1:05 pm PST
Subject: Politics or science

[From Rick Marken (920321)]

Happy birthday dear Bach, Happy birthday dear Bach
Happy birthday dear Johann
Happy birthday dear Bach.

Avery Andrews (920321) writes:

>S-R psychology is a dead-end because it misrepresents what's going on
>in at least three interrelated ways:

Not a dead end -- dead wrong.

>And, just as Gary, Bill et. al. say, even though S->R psychology is
>officially dead, it is still seriously alive as an influence in AI,

It's not only not officially dead, it is a deeply institutionalized component of all life science research. Have you looked at a social science research textbook lately? What do they say is the proper way to do an experiment? Manipulate one or more independent variables and measure their effect on a dependent variable. The assumption is that behavior (the dependent variable) is caused (perhaps indirectly) by the independent variable. This is an SR model no matter what you call the kind of research you are doing -- cognitive, linguistic, AI, mental, schmental (it's cultural) -- whatever. SR is not an influence on AI and other areas of psychology -- it IS AI and those other areas.

>It seems to me that there can be no harm for CT in recognizing this
>kind of thing as an S-R setup,

> Marginality is not the same thing as
>nonexistence, and looking at the properties of real cases where
>S-R setups can work ought to make it even clearer why, most of the time,
>they don't.

One of the problems with PCT is that it is difficult to be politic and right at the same time. As Bill said in his last post, when there is negative feedback from output to input then there is no SR -- that's that. You seem to be suggesting that the politic thing for PCT to do to get accepted by conventional life scientists is to admit that there are some, marginal SR systems. Doing so would be a way of saying "see, we're not saying you folks were completely wrong -- you just didn't notice some of the SR systems where feedback is involved". But there is just no getting around the fact that PCT means that life scientists have been staring negative feedback control in the face for nearly 200 years and seeing it as SR. PCT says that a lot of famous, important, authoritative behavioral scientists are completely and utterly WRONG -- it doesn't say it explicitly, of course, but its pretty obvious once you start looking at the facts and the model of control.

To continue with a tired analogy, PCT is to all conventional life science as the sun centered solar system is to the earth centered one. I imagine it was difficult, at one time, to have proposed the sun-centered model and be seen as anything other than a radical -- an upstart firebrand trying to tear down the well-established foundations of astronomical science. How could such a proposal have been seen as anything else? The same is true of PCT -- there is no way to be a PCTer and not be seen as, well, a lunatic -- claiming that the cherished idols of behavioral science are WRONG. But I don't see any way around it -- that's why so few people are really into PCT. The psychological effects of understanding PCT can be chilling. You get used to it -- but you always feel a bit like Cassandra during the Trojan War.

People will not understand PCT unless they are willing to accept the consequences of that understanding. There is just no way to make the phenomenon and the theory of control more palatable other than by saying that the theory is what it is not. Some people, for some reason, are willing to understand it; others are unwilling to make this radical shift in perspective (and, in so doing, throw out a carefully constructed "baby" that is very precious). I am deeply sympathetic with people who are not willing to give up the the conventional perspective. It is not easy to abandom cherished ideas -- and, frankly, I'm not comfortable encouraging people to do that. All I want to do is try to present PCT as best as I can and hope that others will see some value in it and start working from that perspective.

Best regards   Rick

9203D CSGnet Pg 75 9/24

Date: Sun Mar 22, 1992 8:43 am PST
Subject: Closed Loop Permissions

From Greg Williams (920322)

>Bill Powers (920320.0700)

>Avery Andrews and others in same boat:
>To send copyright permissions for Closed Loop to Greg Williams, you can use
>e-mail. Greg's e-mail address is 4972767@mcimail.com.

Please DON'T! I want paper mail ONLY for permissions, so I can treasure your signatures and sell them to autograph collectors when you all are famous!

Thanks,
Greg

P.S. Why doesn't somebody look at the experiments which have been done on moths escaping from bats? The data might (MIGHT!) forestall a lot of armchair opining. I don't know much about moths, but a LOT of work pointing to open-loop mechanisms has been done on the cockroach escape response (Randy Beer should be well-versed on this data).

Date: Sun Mar 22, 1992 12:10 pm PST
Subject: Open loop behavior

[From Rick Marken (920322)]

I have to put some irrelvant header in here because that's the way my editor at home likes it. So I wish you all a happy and peaceful Sunday.

Greg Williams (920322) says:

>P.S. Why doesn't somebody look at the experiments which have been done on
>moths escaping from bats? The data might (MIGHT!) forestall a lot of armchair
>opining. I don't know much about moths, but a LOT of work pointing to open-
>loop mechanisms has been done on the cockroach escape response (Randy Beer
>should be well-versed on this data).

Excellent suggestion. But I think that the armchair has been under-rated as a tool for scientific research. For example, I would very much like to know about the work pointing to open-loop mechanisms in the cockroach escape response. But I hope the people who did this research sat in their armchairs for a few seconds before diving into the lab. For example, I would like to know what they imagine an open loop system to be. I think of it as one in which some output variable is a function of some input variable: o = f(i). It is open loop if o has no effect on i. A lot of this could probably be determined by anatomical inspection of the organism -- i must be a sensory input and o but be some consequence of efferent neural impulses. A person in an armchair could probably get out the cockroach anatomy text and a physics book and figure out whether the hypothesized sensory input is influenced in any way be the output. If a possible o -- >i influence exists but the researcher still believes that this is an open loop system, then he or she would have to develop a behavioral test to demonstrate that the system actually operates as an open loop system (even though inspection shows that i =g(o)). One obvious one is that the effect of disturbances to o should be exactly what is expected from physical analysis. Have these kinds of tests been done?

One of the big problems in this analysis would be determining the output variable. You mentioned that there have been lots of work on the "escape response". Is this the output variable? If so, how is it measured? It seems like the closed loop goal of the response is incorporated into the verbal definition. So if you look to see if a sensory variable results in that output you are neglecting the fact that your definition of the output includes a closed loop component. That would be cheating. So you must be careful to measure an output variable that does not include the inputs (or disturbances, for that matter) that are also influenced by or that influence the output.

This search for "open loop" systems seems rather ironic, doesn't it? I mean, if the conventional view is correct open loop systems should be all over the place -- it's the closed loop systems that are dismissed as arcane rarities. The fact of the matter is that most of the apparent open loop systems are really just a result of the o = -kd property of closed loop systems. So why all the interest in showing that there are open loop systems? Could it be because people want them to be more prevalent than they actually are (if they exist at all)? Why the resistence to studying closed loop systems the way they should be studied -- using the test for the controlled variable? Why not just let the occasional open loop systems that supposedly exist just pop up out of that research? Doesn't it look a little like the researchers are busy controlling for the perception of open-loop systems -- because that's the way it's supposed to work? Why is the fact that living things are virtually inherently closed loop (because, as I said, their sensors are housed on the devices that produce the outputs) thought of as a minor point -- something like --" sure, I know that outputs affect inputs but that's just adding an extra component to the formula and it's rare anyway". Why don't they want to understand that they are dealing with systems that DO NOT RESPOND TO SENSORY INPUTS. The o=f(i) law is part of a CLOSED LOOP -- i is as much response as stimulus.

Still, I really do want to know about the "open loop" research on the cockroach. I would, frankly, be amazed if there were a convincing demonstration of even one such system in even the simplest organism.

Waiting in my armchair. Rick

Date: Sun Mar 22, 1992 12:45 pm PST
Subject: Re: Open loop behavior

On Rick's view, it doesn't seem to make any difference whether the effects of output on input have any role in the functioning of the device, that is, are part of the story of why it got selected for. As the moth plummets, the quality of the sonar signal changes, but this produces no change in the response. Then the moth hits the ground, and the signal continues, until eventually, of its own accord, the bat moves on, the sonar goes away, and normal activity recommences. (Or is THIS what makes it a control system perhaps? If so, it's hard to see because the moth's output is not causing the relevant change to its input. But maybe that doesn't matter.)

I agree with Rick on the role of armchair debate - the point here is to clarify the conceptual issues involved, so that literature is read and experiments done with smart questions in mind rather than dumb ones.

Avery.Andrews@anu.edu.au

Date: Sun Mar 22, 1992 2:24 pm PST
Subject: Radio Control System

[from Gary Cziko 920322.1550]

I was recently reading Bill Powers "The Cybernetic Revolution in Psychology" in _Living Control Systems_ and took note of his observation:

"Not many [cyberneticists] who led that movement [cybernetics] had ever designed and built a control system, or cursed and sweated to make it work properly, or experienced any extended personal interactions with a working control system; the interactions tended far more to be between cyberneticist and block diagram." (p. 104).

While I don't think that I will ever design a control system or curse to make one work, I couldn't help noticing a "digital proportional radio control system" for $49 in a local hobby shop and so figured that this might be a way for me to at least interact with one (an artificial one, I mean; I already have lots of experience with the living kind).

This is a Hitec "Challenger 260" 2-channel system that includes a pistol grip transmitter (reference level manipulator) and receiver connected to two servomechanisms (it is made for controlling speed and direction of model powered boats and cars). Pulling the trigger and turning the wheel on the transmitter move wheels on the two servos. I replaced the wheels with two four-legged spiders that came with the kit and attached rubber bands to one arm on each.

With either the transmitter or receiver turned off, one can quite move the spiders for a total range of about 90 degrees (it's a bit stiff and I don't know how "good" this is for the servos). But with the both transmitter and receiver on, they really fight to respect their position reference levels. You can feel them vibrate and fight back when you try to disturb then. While it IS possible to overpower them, I am quite impressed at how strong and stubborn the two little servos really are--the more you try to push them around, the more they push right back at you (very much like most people I know!).

The rubber band is a nice way to add disturbances. I can ask someone to pull the rubber band hard any which way and it makes virtually no difference to the position or pattern of movement that I am sending with the transmitter. This is a very nice demonstration of why controlling reference levels is the way to go. I let the servo control system worry about the rubber band disturber and it makes no difference to me, the upper level reference signal supplier.

While Powers's and Marken's computer demos are great, there is something to be said for the real physical interaction that these servos provide. Also an easy way to give my students hand-on artificial control system experience. Highly recommended.--Gary

P.S. My 10-year-old son was also impressed, but made it clear to me that he would be much more impressed if the system were installed inside a racing car or boat! I'll see if I can stall him with a line-tracking robot (follows dark lines on white surfaces by controlling for low infrared reflection) for about $50 that I ordered today from the new Edmund Scientific catalog.

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

Date: Sun Mar 22, 1992 2:24 pm PST
Subject: Seeing the closed loop

[From Rick Marken (920322b)}

Yes, I am doing other things besides posting to the net today. But its raining here in CA (it's been doing a lot of that lately; we need it. Thank god, the greenhouse effect seems to have finally kicked in) so I'm at the computer here -- checking the mail periodically.
--- End of intro bs

Avery Andrews (920322) says:

>On Rick's view, it doesn't seem to make any difference whether the
>effects of output on input have any role in the functioning of
>the device, that is, are part of the story of why it got selected for.
>As the moth plummets, the quality of the sonar signal changes, but
>this produces no change in the response. Then the moth hits the
>ground, and the signal continues, until eventually, of its own accord,
>the bat moves on, the sonar goes away, and normal activity recommences.

Now I think I see the problem. This will have to be based on an armchair guess about how the moth actually works -- but I think the problem is that you are assuming that the input does it's work and then the moth is out of the loop. The problem with the moth example is that the falling seems to be unaffected by the input as soon as it starts. Let's assume that the fall is ballistic -- so that once input (i) reaches some value the fall starts. Now as this uncontrolled falling occurs i is decresing. I suspect that once i reaches some lower level the "curl in a ball movements" stop and the moth proceeds to continue to fly. If the bat is still around the input might start increaing again and the "fall" will occur and end (once i reaches a minimum). The input (i) control is similar to what occurs in a "bang bang" thermostat controller -- the heat is either on or off but the input stays pretty constant (thanks, in part, to the environmental "smoothing" of the temperarure variations resulting from the heater). The moth's falls are probably different durations each time (because the distance it goes to get i to the "threshold" for returning to flying mode depends on disturbances -- wind, updrafts, possibly physical interferences-- so I imagine that what really happens is a series of falls (if the bat remains in the area) of different durations and different distances.

It may be that the fall is always just a one shot thing because it always (or nearly so) gets the moth out of range so that i reaches the minimum threshold to return to flying mode and the moth never runs into the bat again. I don't believe this scenario -- but that's the way it seems from my armchair. But if it does work this way it's still control of input (not open loop) -- and the output (falling/flying) has negative feedback effects on the variable (i) that causes the output (flying/falling). The fact that the feedback effect is negative is seen by inspection --

increase in i lead to increased falling and decreased flying: increase falling leads to decreased i and increase flying CAN lead to increased i. The negative feedback effect may be considered somewhat week (because of the binary nature of the output) but its there.

Does this help?

I bet Bill P. can make my fumbling efforts much clearer. But I try.

Good questions Avery !!!!! Hope things clear up here for you visit to Palo Alto. (I mean the weather, of course).

Regards   Rick

Date: Sun Mar 22, 1992 2:54 pm PST
Subject: Emotions of Memory

[From Kent McClelland 920322]

(Bill Powers 920318.0800)
(Bruce Nevin 920320.0841)

Catching up on the week's mailings on the net, I note that the topic of emotion has come up again. I persist in thinking that the topic may not be getting quite all the attention it deserves. Here are a few more ideas on the subject.

Discussion of emotion on the net so far (e.g., Bill Powers 920203.1000; Kent McClelland 920209; Bill Powers 920210.0900) has focused on the emotions which accompany or result from ongoing actions, i.e., current attempts to control perceptions. I think we may be missing something by not paying more attention to how emotions work in the context of memory. I want to suggest in this post that emotions in conjunction with memories may be important in decision-making and in the sorts of phenomena that non-HPCT folks describe as "motivation."

Let me begin by assuming that when a memories are somehow sedimented in the central nervous system, some record of the accompanying emotions is also laid down, maybe a bit like the sound track of a movie. When a memory reoccurs, then, it often comes with an instantaneous surge of re-experienced emotion. This happens, I think, for prosaic, everyday, habitual sorts of memories, just as much as for recollections of the more traumatic or earthshaking events in one's life.

For example, when I walk to my office in the morning, I'm often worrying about the tasks that lie ahead in an effort to plan out my day. As each remembered duty pops into mind, it is accompanied by a little ping of positive or negative feeling, depending on whether it's something I "feel like" doing that day, often as a result of my emotional state the last time I was working on that task or a similar one. When I have some choice about my priorities for the day, these tiny bursts of emotional memory seem automatically to decide for me the order in which I will tackle things that morning. Things I really want to do (accompanied in memory by that little glint of warm fuzziness) and things I know I absolutely have to do (for which the tingle in my tummy comes from imagining what would happen if I didn't get them done) get the highest priority, while things I can't quite bring myself to cope with just now are put off till another day. In other words, when I have a choice to make, my emotion-tinged memories seem to be making the choice for me.

Now I don't consider myself a particularly irrational or over-emotional person. Quite the contrary, but introspection suggests to me that such emotions often play a part in selecting the memories I use as reference signals for my actions. For instance, how many of the contributors to the net are like me on occasion, in finding themselves unable to refrain from the imagined pleasure of contributing to this conversation, in spite of other obviously more pressing duties?

If decision-making is as emotional as I'm portraying it, what does it mean to be rational? Not, I would suggest, that emotions are absent from a rational person's decision-making, but that the person manages to move "up a level" in the perceptual hierarchy, instead of following his or her first impulse to select the most emotionally attractive option from among the choices that present themselves. If we can assume the point of view of the level of principles or system concepts or whatever, we can weigh possible programs of action by some other standard than whatever feels best at the moment. I'm not arguing that the memories which represent our high-level reference standards are always emotionally neutral, but rather that because high-level concepts are more abstract and thus become averaged out in memory by their application to a variety of situations with different emotional tones, they may end up being less emotionally charged than lower-level memories. Low-level perceptions, say things on the configuration level, like colors (one's favorite color), odors (fresh-baked bread), or sounds (the rat-tat-tat of a jackhammer), often carry strong emotional associations. [The sociologist in me wants to point out that ritual occasions may be intended in part to teach people to associate the proper emotions with their system concepts, as when we all sing the national anthem in an effort to feel patriotic.]

Once we have made a decision to do something, the emotional memory attached to the reference standard defining our goal may then help us to persist in that line of action till the goal is finally achieved. Complicated programs of action often take a long time to unfold, and for much of that time feedback from the environment may be sparse. Our own attempts to stay on track toward the goal may be intermittent, as we pursue other aspects of our life at the same time. What can keep a person slogging away at a distant objective, when the immediate feedback is mostly discouraging or nonexistent? Well, it seems pretty obvious to me that if we imagine the successful completion of the project and anticipate in imagination the accompanying glow of approval and satisfaction, it will help to "motivate" us against present adversity. Depending on our level of "emotional maturity," this fantasy fulfillment need not be an elaborate daydream. Just a little ping of positive feeling might be enough to to get us back to work.

I realize I'm beginning to sound like Carver and Scheier here, as they babble on about "expectations," or worse, like a Freudian delving for the emotional wellsprings of behavior, or even (shudder) like an S-O-R psychologist proclaiming the importance of "self-reinforcement." But I think that all these people in their various misguided ways have been on to something which we might be able to describe more clearly in the HPCT framework. Maybe my student who made the ostensibly absurd suggestion that emotions were a kind of "12th level" of the hierarchy wasn't so very far off the mark (see Powers 920210.0900). Maybe the emotions associated with memories can give us a key to "where the 'wants' come from."

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

Date: Sun Mar 22, 1992 3:41 pm PST
Subject: Re: Seeing the closed loop

I'm happy to see the hypothetical moth *you're* describing as a control system (more than happy - that's obviously what it is) but it's *not* the one *I'm* describing, whose plummet causes it to crash into the leaf litter and lie there until the bat goes away.

The reason I'm being persistent in the hunt for S-R systems is not to uphold the reputations of the founder's of modern psychology (these guys mean nothing to me - I'm a Chomskyan linguist), but due to a liking for the singular and the bizarre, and a distrust of claims for uniformity.

Avery.Andrews@anu.edu.au

Date: Sun Mar 22, 1992 5:20 pm PST
Subject: Open loop/Evolution

[Rick Marken (920322c)]

a
b
c

> it's *not*
>the one *I'm* describing, whose plummet causes it to crash into the
>leaf litter and lie there until the bat goes away.

Ah, so that's what happens. OK. How does it know that the
bat went away? Still looks closed loop to me - just a pretty
boring loop, if that's all that's involved. The thing remains
down until the input is low for a certain length of time.
Or does the moth always remain in the litter for exactly the
same amount of time when it falls?

Do you know so more real details about this moth? It does
sound very intteresting.

Regards   Rick

Date: Sun Mar 22, 1992 6:14 pm PST
Subject: S-R, CT, Evolution, and stuff

[From Bill Powers (920322.1700)]

Rick Marken called today and we commiserated for a while about an odd phenomenon on the net: we're down the the nitty-gritty difference between S-R theory and control theory, and S-R theory is getting a staunch defense! This is actually a good sign: if S-R theory weren't so seductive, ten generations of smart people wouldn't have fallen for it. Control theory is not up against a trivial opponent.

I'll ask this again: isn't it odd that the same stimulus that causes a reponse seems so often (always?) to be affected by the response? You wouldn't think that evolution would attach this condition to responses -- unless there were some very strong survival advantage in acquiring responses that control inputs rather than responses that simply alter an objective effect on the physiology of the organism.

The answers are found in control theory. A stimulus, in the final analysis, is all the organism knows about its environment. By "stimulus" we don't mean the distal or remote event that the casual observer sees in the environment, but the actual energy impinging on the sensory receptors of the organism, wherever the energy came from. If the nervous system is going to do anything behaviorally appropriate to survival, it must base its actions not on what's "really" going on in the world, but strictly on the signals arising from its receptors.

And how is the organism to know that it has made a move that is actually appropriate to something important outside it? Only through the effects of the action on the sensory signals. If the action has no effect on a sensory signal classified either as dangerous or as beneficial, then the action is probably useless. If an action doesn't decrease the "dangerous" signal or increase (to the approriate level) the "beneficial" signal, the action is just a waste of energy.

The judgment of dangerous or beneficial can be made in many ways: it can be made rationally, it can be taught, it can arise from internal reorganization, or it can be effectively made by natural selection. The whole point of control theory is that whichever process is at work, the result is a control system, not an S-R system. The result is a system of the kind that acts to modify the very input information on which the action is based. And to be more specific, the effect of the action on that input is in the direction that brings that input closer to the state that is better for the organism. The action might affect the cause of that input, or the organism's relationship to that cause, or even just the input signal. The path by which the action has its effect is irrelevant: all that needs to be known is that the input correlates with survival, positively or negatively, and that certain actions can affect it.

Does this mean that control theory denies natural selection? Not at all. The reason that little Soldier Beetle I experimented with is colored like bark is that it spends a lot on time on bark or on the ground. It's hard to see in a natural habitat. The variants that are easier to see get eaten. Viola: we are left with the hard-to-see ones. But I don't think this concept is enough to account for behavioral organization in general.

We, as sophisticated human observers, can see that the stimulus in itself isn't better or worse for the organism -- it simply stands for some external situation that is better or worse in terms of actual effects on the organism. These actual effects might be effects on hunger or thirst or overheating, or they might be effects on ability to reproduce. The mechanism that sees to it that the action reduces dangerous effects or increases beneficial ones might be a learned control system culturally transmitted, a reorganization that enables a baby to get food when it needs it from the particular caretakers at hand, or Darwin's Hammer that simply removes the organisms that guessed wrong. But in all cases, the driving force behind acquisition of these peculiarly appropriate behaviors is inside the organism, in its own capacities to change (by swapping genes, by producing random changes in organization, or by memorizing new reference signals). And in every case, the result is a control system, not just a passive arbitrary response to stimulation.

There are some holes in a simplistic S-R/evolutionary concept of how behavior gets to be organized as it is. The basic story is that inputs get hooked up to outputs, in simple organisms, by natural selection. But there are all kinds of annoying little phenomena that don't jibe with this view. I suspect that if biologists were actively looking for such deviant phenomena, they would find an abundance of them.

Consider what Beer said about waxing the cerci of the cockroach. This throws off the detection of the direction of winds, so the cockroach's escape response goes in a wrong direction. That's what you'd expect if the escape response were wired in, being changeable only from one generation to the next. But Beer then said that if the threatening puff of air is repeated again and again, something changes inside the cockroach, the very same live cockroach within its own lifespan, such that the escape response becomes appropriate again. Obviously, evolution had nothing to do with this change, because cockroaches have to die before evolution comes into play. The internal machinery of the cockroach somehow detected the inappropriateness of the escape response, and altered itself to make it appropriate again.

What criterion could have directed this sort of change? The only possible one is that the escape response was not having an opposing effect on the "bad" stimuli that gave rise to it, over many trials. The change of organization stopped only when the escape response was in the direction that once again took it downwind, toward smaller velocity fields, and thus led to a lessening of the stimulation.

If information from the cerci can be reconnected so as to bring cercus signals back under control after external interference, why couldn't the original organization have arisen the same way as the cockroach matured? Why couldn't there have been some roughly correct initial organization that then became reorganized to produce the adult "reflex" we see? Why couldn't the effect of evolution be simply to define certain input signals as "BAD?" If a reorganizing capacity exists, and if it worked effectively, that would be enough to result in acquiring hookups that used existing behavioral output systems to control the bad signals toward the evolution-selected reference level of zero. It wouldn't be neccessary for evolution to anticipate the physical effects of the approach of large predators in detail. The sensory signal tells enough of the story, when compared with the reference level, to go on with.

There has always been a bias in biology and related sciences to see organisms as machines pushed around by the environment, and even selected by the environment. It has seemed important to explain everything about organisms without going beyond the principles of physics and chemistry: it's a clockwork universe, and everything happens as it was forordained to happen when the Big Bang brought physics into being. One doesn't have to pursue this line long before the motive becomes clear: specifically, living systems are natural products and not supernatural ones. We don't explain behavior, in science, by referring to souls or consciousness or -- the big one -- God's Purposes. If you start to talk to a biologist about the purposes residing inside an organism, it won't take five seconds before you start seeing sideways looks and nudges and winks being exchanged in the audience. You might as well have started telling a funny story about a nigger, a kike, and the Pope. Speaking of purpose is, scientifically speaking, in bad taste or worse.

The result is that evidence of purposiveness is systematically suppressed, ignored, and distorted by biologists and their descendants -- or has been, until very recent times. This means that even naturalistic observations carry a strong bias: one doesn't hear about what a moth is trying to accomplish by its behavior, even though the intent may be perfectly obvious. Not just biologists, but most people, interpret words like "purpose" and "intent" to mean conscious purpose and conscious intent, and even more specifically, purposes and intents that can be described in words. Biologists have gone futher; seeing that "purpose" refers to some metaphysical human experience, they have pre-empted the term to mean "consequence that affects fitness."

Clearly we don't expect conscious purpose to exist in the world of a moth. The idea that purpose and intent mean something far more basic than our consciousnesses or our descriptions would simply never occur to anyone -- who was unaware of control theory.

I think it is quite reasonable and proper to enquire about the role of purpose or intent in evolution. But this will be almost impossible to explain to a biologist or geneticist, because it will seem to mean outside direction of evolution, and nothing said after that will be heard. But if one doesn't have this allergic reaction, the idea can come to make a great deal of sense. To make sense of it, however, you have to stop thinking in terms of human intentions and purposes, and boil the concept down to its essential concept: control of input relative to a reference level.

The "action" of a system is its effect on the local environment, including its physical self. The "sensory input" is the effect of the local environment on internal variables that are sensitive to external energy impingments, such as chemical interactions. At this basic level, even DNA can act on its environment and sense the state of its environment. There is no reason why DNA could not contain, or regularly manufacture, comparators for judging deviations of certain variables from reference levels defined in the code. There is no reason why this sort of control organization could not have developed a number of hierarchical levels, even in so small a world as cellular genetic chemistry.

The most basic level would be a reorganizing system. This system would monitor certain essential variables, those which sensitively indicate the viability of the genetic structure, whatever that might mean. When external conditions change in such a way that their effects can reach even to this level, overwhelming all other control structures erected to protect these basic variables from disturbances, the reorganizing system will start inducing random changes in the very code of the DNA. It will continue to induce random changes until the essential variables are protected from further disturbance.

Clearly, this reorganizing system can't have any inheritable effect in the lifetime of a single organism prior to reproduction. But much more is passed from parent to offspring than just DNA. Even in the human being, for example, the mother's mitochondria and other intracellular structures are passed along, and this means that control machinery and controlled variables can survive from one generation to the next and keep right on operating. So this basic control system can work across generations, the coming and going of individual hosts being just a brief bobble that occurs now and then.

Without elaborating this idea further, it should be plain that control theory offers the possibility that purposes may be behind evolution itself. While evolution offers no guarantees of success, and is not directed from outside in the slightest (certainly not by the environment), it is obviously directed toward the organism's gaining greater and greater control over what happens to it. The direction comes from inside the organism, from its genetic machinery. Of course because timed reorganization -- blind variation -- is the engine behind this building of greater and greater control, the variations that occur are quite random, or as random as gene-swapping plus truly random change allows. Variations that render the organism unviable will be wiped out. Organisms that survive more surely to the age of reproduction will predominate over those that don't. The basic arithmetic of natural selection continues to apply.

But there is, nonetheless, direction: the organism reorganizes faster under stress, slower when it is successful in controlling the variables on which accuracy of replication depends. That's the logic of reorganization. This inner purpose sees to it that evolution progresses as well as moving sideways and backward. The measure of progression is not complexity or size or even numbers, but the capacity to control. Even within a single species, the capacity to control can keep improving generation by generation; today's cockroaches, which look just like those of 300 millions years ago, may contain far more sophisticated control systems than those first ancestors had. Today's human beings certainly have vastly more control over what the environment does to them than did their ancestors even 10,000 years ago. "Selection pressures" of 10,000 years ago are now effects we produce for fun.

I have brought all this up to show that the CT viewpoint has to be carried all the way through, as far as older concepts have been carried, before its worth can truly be evaluated. The arguments that have been going on on this net have all been superficial; the arguments for natural selection of behavior as the only important force rest on layers and layers of assumptions about the nature of organisms at all levels to the genetic level, and all those supporting layers are permeated with S-R assumptions. The surface thinking may change toward a CT point of view, but unless the supporting layers of reasoning and assumption also change, to encompass the same principles, the surface change can't survive. CT gives us a new perspective not just on overt behavior, but on the functioning of organ systems, hormone systems, cells, genetic processes, and basic biochemistry. It is simply a new point of view toward the processes of life.

Finally, one last comment. In Avery Andrews' last post (920322) he said

>Then the moth hits the ground, and the signal continues, until
>eventually, of its own accord, the bat moves on, the sonar goes away, >and
normal activity recommences.

This is a beatiful picture of what is basically wrong with S-R theory EVEN IF THERE ARE S-R RESPONSES. After the big dramatic event is over, the observer turns away, because now "normal activity recommences," and in normal activity there is almost nothing that can be understood from the S-R perspective. It is typical of psychology that finding responses to stimuli is very difficult -- a graduate student can try out a dozen possibilities before a coherent effect pokes its head above the surface of noise. S-R theory works so badly that 99.999% of the action that goes on in an ordinary organism pursuing its ordinary activities is unexplainable. Most of the time, organisms are "between responses." Even when some action sticks out enough to be called a "response," there's usually no hint of the stimulus -- the term "response" is a religious term.

Under control theory you can close your eyes in the middle of a crowd, spin around three times, open your eyes, and see somebody controlling something, a dozen things. The flood of examples is overwhelming. Every action is resisting some kind of disturbance of something, if only gravity. Every act varies until some specific result occurs. You can test a hundred control systems for resistance to many kinds of disturbances in one trip across the meeting hall to the drinking fountain, and you might even notice what you're controlling for yourself.

It's laughable that we should be wasting our time debating about unusual phenomena that stand out precisely because they're different from almost all normal behavior. It's normal behavior, usual behavior, everyday behavior that we have to account for first, before we can even tell what is unusual behavior.

Best,   Bill P.

Date: Sun Mar 22, 1992 6:21 pm PST
Subject: betray

Well, I gave pretty lousy semantics for betray. Here's a second attempt (still sloppy, I'm already spending more time on this that

I ought to be).

```
  X betrays_1 Y: X expects Y to do things that are good for X.
                 But, wanting bad things to happen to Y,
                   X does something which lets people who want
                   to do bad things to Y do them.
```

This is the core sense, the one conveyed also be the noun `treachery' (if my sleazy appearance and shifty gaze betray my intent to defraud you, this isn't treachery. Neither, I think, is cheating on your spouse, though I wouldn't be too surprised if people differed on that one).

There will be a various other senses besides this one, a bigger list than people might be initially comfortable with, but not infinite (I'm taking there to be a difference between creative metaphors and the ones that have become standardized as additional senses). Here's another one:

```
  X betrays_2 Y: Y is something about Z that Z does not want people
                   to know about.
                 Because of X, people come to know about Y.
```

e.g.
  Jack's accent betrayed his humble origins.
  The Playboys under the mattress betrayed John's poor taste in porn.

And a (the?) third:

```
  X betrays_3 Y : Y expects X to be faithful to Y.
                  X doesn't be faithful to Y.
```

I would claim that if you can't understand the notion of sexual fidelity, you can't understand betray_3, though I don't have an explication for this notion at the moment. Otherwise, these are pretty close to Wierzbickian primitives, though the syntax still needs a lot of work.

Fire Away!   Avery.Andrews@anu.edu.au

Date: Sun Mar 22, 1992 6:30 pm PST
Subject: Re: S-R like control theory???

[Martin Taylor 920322 2120]
(Bill Powers 920320.1800)

I'm having problems with bad lines and inserted characters again, so be tolerant of errors, please...

I agree entirely with:
>
>S-R theory is a subset of control theory only in the sense that the control
>model explains why particular stimulus-response laws are observed (when
>control is in fact present). At the same time, control theory shows that
>these laws are illusory; they describe not the organism but its
>environment.

>
>So there is very little in the S-R behavioral literature that actually
>increases our knowledge about the properties of organisms; S-R formulations
>are not just somewhat simpler or alternative descriptions of the same thing
>that control theory describes. They are, when control is actually present,
>misrepresentations of the organization of behavior.
>
I'm on your side, you know! I was pointing out that when a reference level is fixed and you apply a disturbance that affects the percept, you get an error that causes a more or less reproducible behaviour. That looks like S-R effects if you are predisposed to look at things that way. So if, as you previously suggested, we do experiments "simply", by fixing reference levels and looking at things like tracking, then we will have problems distinguishing the results from those of S-R theorists.

As far as I can see, PCT can predict anything S-R can, but the reverse is not true. In that sense, S-R is a subset of PCT. It works under very restricted conditions, or at least it appears to, even when there is real control.

Actually, on reflection, I disagree with "control theory SHOWS that these laws are illusory." No theory can show that another is an illusion. Belief in one theory can lead to a belief that another is illusory, and that belief is easily generated when one can explain all that the other does, and then some. But you can't get away with the quoted phrase if you are talking with someone who is agnostic or of the other faith.

All the same, data obtained by people who believe in S-R can be useful to PCT theorists, if we acknowledge the relation among reference, perception, and error in the feedback loop.

Martin

Date: Sun Mar 22, 1992 6:47 pm PST
Subject: open loops, & Open Loop

[Avery.andrews 920322.1144)]
 (Rick Matkrn (920322c))

>Ah, so that's what happens. OK. How does it know that the
>bat went away? Still looks closed loop to me - just a pretty
>boring loop, if that's all that's involved.

So, the loop is closed even though R exerts no influence on S. OK, if you want to call it that way, but there might be something to say for distinguishing this kind of boring loop from the more usual and interesting ones.

No more details about this critter, since it's hypothetical, but here's another one that isn't: male erections, as a response to various kinds of stimuli, including pornography, etc. Viewed in a wider context, this is presumably part of various closed-loop systems, but on its own, it still looks pretty open to me.

Note to Greg Williams: my permission went into the mail a few days

ago, so if you don't have it in a few days, holler.

Avery.Andrews@anu.edu.au

Date: Sun Mar 22, 1992 7:06 pm PST
Subject: "Ballistic"

From Greg Williams (920322)

Maybe a better word than "open-loop" (which was an attempted replacement for "feedforward") is "ballistic." As the gun is pointed and fired, so the cockroach seems to go off on its escape trajectory, not controlling against disturbances on the way, landing it cares not where (with regard to the air-puff perceptions which "pulled its trigger"). Or so says Camhi and colleagues (his textbook provides a good review of the neuroethological experiments). Certainly, the cockroach's ballistic trajectory results from control for the perception of no-air-puff. But the operation of that control is ballistic, not continually corrected during the trajectory to, say, attempt to keep heading in a direction with less air-puffing. Rather similar to the non-continuous control exhibited by E. coli for getting food, I think. So what is the objection to having ballistic mechanisms and non-continuous control in PCT?

Greg

Date: Mon Mar 23, 1992 8:48 am PST
Subject: input signals

This message is to Bill P.
Please send replies to my personal address. I would send this directly to you but I cannot make out the address on the screen--you seem to have some unusual symbols in your address.

Anyway, I haven't been keeping up with the net lately so if this is being discussed now, forgive my negligence. I'm having difficulty understanding the quality of the input signals for higher level systems. I know that an input signal for one system is an integration of (   ) from lower systems--what's in the (   )? Are the signals that become integrated simply the same as the signals the first level receives? Or are these signals somehow adapted? All the diagrams I've seen make it seem as if they are not adapted, tht they simply go all the way up. But if they are integrated then something is different. I'm close but I don't quite have it.

On another note, I agree with your evaluation of the kid and the fire alarm, that it was probably luck. I still contend, however, that other ways of knowing should not be deemed impossible. I prefer the luck explanation, myself, but I KNOW that there are instances of knowing which cannot be explained by luck or traditional means.

To Mary,
Did you receive my message about a month ago about Closed Loop copies. I have friends who would find our discussions very interesting so I would like to get them copies of Closed Loop, including the last one (social control) and the previous one (higher system concepts). Should I send you $10 with their addresses? Can they get back copies like this?

Carpe' Diem
Mark Olson (m-olson@uiuc.edu)

Date: Mon Mar 23, 1992 10:51 am PST
Subject: open loops, & Open Loop

[From Rick Marken (920323 8:30) I guess if I'm going to post this much I ought to start time stamping]

Avery Andrews (920322.1144) says:

>So, the loop is closed even though R exerts no influence on S.

No, it's closed because the fall (R) does influence S (the sensory input that led to the fall).

Greg Williams (920322) says:

>Maybe a better word than "open-loop" (which was an attempted replacement
>for "feedforward") is "ballistic."

Ok. The fall itself is ballistic. I agree that there are often ballistic components in closed loops. In fact, I think there might be a case of an open loop (in my sense -- o = f(i) but i<>f(o)) in the moth fall. What the sensed sound actually causes most directly, I believe, is the wing folding (well, contration of the muscles that brings the wings into a ball). So m = f(i) where m is muscle tension. The muscle tension causes the wing collapse (w) so w = g(m). Finally, acceleration is a function of the degree of wing collapse a=h(w). I think it is possible that the relationship betweem a and w could be open-loop -- or, at least, if closed, the loop gain could be relatively low. As the bug accelerates dowm there might be some turbulance acting to open the wings -- this would be negative feedback. But if the muscle tension tends to bring the wings into a closed circular shape I could imagine that the effect of the wind generated by acceleration could be negigible (open loop) or could act to keep the wings closed (positive feedback).

Ultimately, it depends on what you mean by "behavior". At the highest level this moth is controlling it's perception of sensed sound (probably) -- and if this were seen as a response to stimulation this would be a mistake. This mistake could be prevented by realizing that the behavior of this organism could be involved in control of a sensory variable or variables, hypothesizing what that variable might be, doing "the test" and revising the hypotesis as necessary. This approach to studying behavior would reveal any variables that were uncontrolled (open loop) but that could be shown to be part of the physical means used to control an input. When open loop components of control loops are found you probably would be reluctant to call them "behaviors" of the organisms -- since they are not intentionally produced. But that's a matter of taste more than science. If acceleration is an open-loop consequence of the "shape" of the organism then that would be an interesting finding -- and, in fact, one that could ONLY be revealed by the test for control.

Which brings me back to my question from a couple days ago: how do researchers know about the open-loop organization of the cockroach escape response? I have a feeling it is based on the wrong data -- but that does not mean that there are not, possibly, some open-loop components of this behavior. I just can't understand how any "behavior" that consistently achieves some result (escape) could be completely open loop (unless it were always performed in the same environment, from the same orientation, with the same motor charateristics, etc.). The achievment of consistent results in a variable environment is control, by definition. And the only way we know know how to make systems that control is to build systems that sense and have negative feedback effects on the result that is to be produced consistently.

Regards Rick

Date: Mon Mar 23, 1992 1:37 pm PST
Subject: Misc Remarks

[From Rick Marken (920323 12:30)]

Bill Powers (920322) -- you mentioned some comments by Beer about wax on cercae? I don't remember anything like that in Beer's posts. Have you heard from him lately by the way? How's the bug comin' along?

Kent -- Your post on emotion was very interesting. I'll try to respond ASAP. Otherwise, we can discuss it in person in KC.

Gary -- Thanks for the info on the controllers. Yes, it is a good idea to interact with "real" servos. Something like what you describe was built by Bill P. as his "arm" demo -- not a computer model. A real position control servo with motor output that kept an arm pointing in the direction to which you set the reference. I mention this only to show how hard it is to think of things that Bill hasn't thought of already. But your controller seems more portable. I think all PCTers should have one -- along with the rubber bands and coins.

Hasta Luego Rick

Date: Mon Mar 23, 1992 2:51 pm PST
Subject: Re: open loops, & Open Loop

[Martin Taylor 920323 16:30]
(Rick Marken 920323 8:30)

Rick, I think your posting on the moth and the larger question of behaviour and open loops brings us close to agreement. There COULD be open loop behaviours in which the loop is closed by evolution, but it seems unlikely that such "behaviours" would survive a long evolutionary path, because circumstances DO change, and actions that are effective in one situation could be deadly in another.

But, in the field of HCI, we see numerous situations in which the interface designer has precluded the closing of the loop, at least forced it to a higher level. The user imagines the state of the computer, and issues a command intended to affect that imagined state. The designer may not have provided any way for the user to determine either the new state or the effect of the command as a state difference, unless the initial state is sufficiently wrongly imagined that the command is a "syntax error." I think one of the main benefits of direct manipulation interfaces is that they provide continuous possibilities for the loop to be closed. The user can see all the time how the actions are affecting the perceived state of the machine. This aspect of the benefits of direct manipulation is never (I think) discussed. Instead, the idea is put about that such interfaces eliminate the possibility of error, which they don't.

This may be my last posting from this address. Gary has moved my reception of the CSG-L postings to another machine, at my request. I will henceforth be found at mmt@ben.dciem.dnd.ca

Martin

Date: Mon Mar 23, 1992 3:54 pm PST
Subject: "completely" closed loop = ?

From Greg Williams (920323)

>Rick Marken (920323 8:30)

>Which brings me back to my question from a couple days ago: how do
>researchers know about the open-loop organization of the cockroach
>escape response?

I've learned my lesson regarding posting synopses of data. You'll have to read it yourself. I suggest that you begin with J.M. Camhi, NEUROETHOLOGY, 1984.

>I just can't understand how any "behavior" that consistently
>achieves some result (escape) could be completely open loop (unless it were
>always performed in the same environment, from the same orientation, with the
>same motor charateristics, etc.). The achievment of consistent results
>in a variable environment is control, by definition. And the only way
>we know know how to make systems that control is to build systems that
>sense and have negative feedback effects on the result that is to be
>produced consistently.

I'm not sure I know what you intend by "completely open loop," but I think that the cockroach's controlling for perceiving no-air-puff can be achieved quite reliably via the following control mechanism: air-puff coming from a certain direction affects sensory hairs in particular ways, resulting in precalculated patterns of stimulation to leg muscles, usually (but NOT ALWAYS, since the "calibrated" action is NOT protected against disturbances) moving the bug in a direction away from the air-puff source. I suspect that highly consistent (even in the face of disturbances), "completely closed loop" (?) escape actions would be less economical in this situation, and/or might be slower than precalculated, calibrated action. Ongoing, fine control of the bug's trajectory simply isn't necessary to (almost always) escape. Still, there is a closed loop here, even if (again, guessing at your meaning) it isn't "complete." In a sense, it IS complete, as the loop encompasses the ballistic movement to get "away" -- the bug (generally) does get away, thus achieving control. BUT ONGOING CONTINUOUS CONTROL ISN'T REQUIRED. One might say that the control is continuous, since the sensors are continually receiving information about air-puffs, but I claim it isn't continuous, because no use is made of that information except to "aim" and "fire" each "triggered" ballistic action following an air-puff.

Greg

Date: Mon Mar 23, 1992 4:22 pm PST
Subject: Purpose, organization, "edge of chaos"

[From: Jeff Dooley 920323.1100]

(Bill Powers 920322.1700)

Since becoming a subscriber to this net some months ago I have learned a great deal, and many concepts of control theory which had eluded me in the past have become, if not clear, at least (apparently) less confusing. For having facilitated that learning I want to thank you all.

I'd like to comment briefly on two issues that popped out at me from Bill's post: (1) the scientific/philosophical difficulties surrounding the concept of "purpose," and (2) the phylogenetic transmission of "control machinery." I'm beginning to think of these as potentially related issues.

Talk of "purpose" is routinely dismissed among scientists and philosophers, as Bill suggests in his post and in his article "On Purpose" from _Living Control Systems_. But why does the concept have such a bad reputation? Further, the question arises, can the notion of "purpose" be decontaminated for sci/phi discourse?

A possible response to the first question centers on the general anti-intentionality of logical empiricism and, later, of philosophical behaviorism. These traditions seemed to assume that "purpose" was cognitively meaningless since it was assumed that such a property was intentional--that there was no possible observable phenomenon to which it could be connected. This appears to be a fundamental assumption which PCT calls into question. According to PCT, as I understand it, purpose may be inferred, inductively, as the maintenance a controlled phenomenal variable at or near reference. Find a phenomenal variable apparently controlled under perturbation and you have located evidence of purpose. As a response to the second question, could not the PCT view of purpose as a matter of observation in this sense help decontaminate the concept for philosophy and science?

Meanwhile, what's worse for the notion of "purpose," it has got entangled with the Aristotelian concept of teleology, by which, as Ernest Nagel has derisively quipped, future events are seen as agents of their own realization. To linear-thinking, cause-effect, (S-R?) philosophers of science, this sounds like laughable illogic. While under a linear model of causality (development, progression, etc.) such a notion as "purpose" appears a paradigm example of bad science, reconsidering it under a circular, equilibrating, control model may cast the whole situation in a new and possibly fruitful light. This brings me to my second issue: the idea of the phylogenetic transmission of the reorganizing system, and to how I see these two issues as related.

I'd like to suggest the idea, blending concepts of control theory purposiveness with a thesis of biologist Stuart Kauffman's: that species (complex, adaptive systems) evolve (drive themselves) to the edge of chaos and maintain themselves there. For Kauffman, "the edge of chaos" is the phase-transition space at which order melts into chaos. Too much order and the system can't adapt to perturbations; too much chaos and development can't get a foothold from which to progress. If we postulate an "organizing function" of a species (or any cocontrolled variables.
Could not a homeostasis at the edge of chaos be considered a controlled variable? Some current work in biology does consider the idea of trans-individual organizing functions working on the phylogenetic level, (in fact, on all hierarchical levels from RNA to species!), but they have not, to my understanding (not being a biologist), figured out a way of modeling the dynamics of progression.
It is precisely in the reintroduction (decontamination) of the concept of purpose in scientific discourse, under PCT, that we might begin to progress in our learning about the possible organization functions of complex adaptive systems, on whatever recursions of organization they become apparent.

Finally, a footnote on different ways of "seeing" the world: The two paradigms of S-R, open-loop, etc, and PCT, closed loop, etc. seem incommensurable in a Kuhnian sense. Workers in the different traditions really do seem to think, speak, operate in, or at least "see," different worlds. This could be serious, since theoretical development, experimental design, and other activities influencing the direction of science are informed, one could say, as a function of which world one belonged to. Could incommensurability in this sense help explain the apparent fact that S-R types and PCTers seem to just talk past one another?

jeff dooley
dooley@well.sf.ca.us

Date: Tue Mar 24, 1992 2:51 am PST
Subject: Levels of perception
Bill Powers (920324.0300)]

Copy direct to Mark Olsen, plus CSGnet

My address is powers_d%flc@vaxf.colorado.edu. Between the powers and the d is an underscore (shift hyphen) and between the d and the flc is a percent sign.

Send Closed Loop orders to Mary Powers}?, $5 each copy.
CSG
73 Ridge Place, CR 510
Durango, CO 81301

Mark Olsen (920323) --

You ask about the functions relating one level of perception to another. This is indeed the question that HPCT poses -- but doesn't answer. What lies behind HPCT is not any proposal as to how each level of perception is derived from the one below it, but a proposal as to what the levels of perception are and how they are related. This is the phenomenon that any model must in the end explain.

The "H" part of HPCT can be taken in two ways: first, as a general sketch of a hierarchy of control in the abstract, with the communication between levels consisting of a series of perceptual re-representations of reality and a corresponding set of reference signals used to control lower levels; second, as a series of proposed levels of perception (and control) based directly on an analysis of experience with the hierarchical control concept as a guide. This is a beginning model; there may well be other modes of communication between levels, but the basic one is probably valid.

The definitions of levels define the modeling problem. We can see that the sensation level is probably derived by weighted summations of intensity signals, the weights defining a vector in a perceptual space having fewer dimensions than there are different sources of intensity signals. But that answer to the modeling problem comes after noticing that sensations seem to depend on intensities in a particular way, a way that could be modeled as weighted summation. The phenomenon to be modeled comes before the model.

And that's as far as I can go. I don't know how configurations are derived from sensations -- how it is that we can get the sense of, say, a particular person's face over a range of distances and orientations and expressions. If signals standing for the dimensions of a face existed, then it's possible to make a rough guess that transitions of the face from one state to another would be sensed using time functions and partial derivatives; that's a feeble start toward a functional model that you could run on a computer. As to the rest of the levels, the kinds of computations involved are mostly a mystery to me. The few guesses we have come up with are strictly stabs in the dark. You can use words like "integration" to describe how some kinds of perceptions are put together to create others, but the word is just a noise. It doesn't tell us anything about the processes involved.

Behind this exploration of perception lies a fundamental postulate; if you don't internalize it, I don't think you can even get started on the problem of modeling the brain's perceptual systems, or for that matter, in understanding HPCT. The postulate, simply put, is this: it's all perception.

By that I mean that no matter what you attend to in the world of experience, whether you refer to inner or outer experiences, concrete or abstract, verbal or nonverbal, the object of your attention is a perception. You are looking at or otherwise experiencing the brain's perceptual activities, not the objective world itself.

Vision is the most important sense to understand this way if you're sighted; understand vision and the rest (touch, taste, sound, etc.) will follow. The world you see begins as pixels (individual picture elements). The pixels are so close together that you see no spaces between them, although the sensory nerves do not overlap and in fact do not completely fill the retina. There's a world between the pixels, but we don't see it unless the view shifts slightly -- and then what we had been seeing disappears into the cracks between the pixels. This is invisible to direct experience; the world seems continuous over the whole visual field. We get a sense of seeing the world at infinite resolution, and can't imagine what the whole field would look like if we had, say, ten times as many retinal receptors and the optical acuity and brain power to take advantage of them. This would be like seeing the world through a magnifying lens, except that the whole world would look that way, not just one little part of it (which we still see at human resolution). The only way to imagine this is to go the other way: view the world at a lower resolution, as in a halftone photograph or a television screen seen close up, and imagine that the result is the only world you can ever see. That's how our picture of the world would look to a different organism with higher visual resolution. But we experience it as having continuous detail right down to the level where it appears smooth. I suppose the fly sees the world in the same way. But its world is smoother than ours.

Building up definitions of the rest of the levels in the hierarchy is then a matter of noticing persistent types of structure in this world of picture elements. The first level above the pixels themselves is sensation, a type of perception that can't be analyzed in any way except into variations of intensity. Color is a sensation, as is shading.

Perhaps things like edges are sensations, derived in one step from the pixel distributions. When analyzing perceptions, however, don't use any data but your own experience. Theory and neural data will tell you that in the visual field, in the retina itself, all edges are enhanced, so that there is a strong outlining effect. But look at the edge of a sheet of paper on a dark tabletop. There is no outline. The closer you look at the edge, the more nearly it seems to be an infinitely sharp line separating uniform white from uniform dark. The edge itself is there -- but you can't see it as an object. It's just a sense of edgeness. Only under special conditions, as in looking at a smooth gradient of illumination going over a relatively short distance from white to black do you see edge effects like the "Mach band", the only clear subjective evidence of edge enhancement. However those neural signals enhanced at edges are processed, the result is that step changes look like step changes, not outlines as in cartoons. Whatever model we come up with for how the nervous system processes pixel information, it must result in edges that look this way, without borders. If it doesn't, the model is wrong.

The next step is to notice that the edges and corners and broad white areas of the piece of paper add up to -- a piece of paper. If you've made this transition properly, it will come as a surprise. Where did that piece of paper, or piece-of-paperness, come from? It wasn't there in the edge, or the corner, or the whiteness, or the darkness. It comes into being only when all those elements are seen grouped into a thing, a configuration with a familiar shape, orientation, distance, size, and so on. The Gestalt psychologists of old spent a lot of time looking at things like these. They should have kept going. Or perhaps they shouldn't have been cowed by the behaviorists.

You have to go slowly and by the smallest steps you can devise. If you go too fast you'll miss the smallest steps; if you miss the smallest steps you'll lose the sense of examining perceptions, and start projecting the visual field into an external world again. You'll jump to the more abstract levels and lose the connection from one level to the next. This is, if you like, a form of meditation on experience in which you distance yourself from experience and look at it merely as a display. You're not trying to see anything about the world, but only something about the display. You're trying to see what features the person who constructed it thought of putting into it, just as when you read a program you think to yourself "Now he's setting up an array to hold the results" instead of just reading the code, or when you read a novel as a literary critic you think "Now he's introducing tension" instead of just getting tense. Who the "he" is is immaterial -- the point is to see what is before you as a construction that has inner organization, and try to see how it is put together.

The general principle is that when you have found a level, like sensation, the next level is going to depend on it; also, the current level depends on the one below it. If you analyze a perception to see what it is made of, at first you see just more perceptions of the same level -- big configurations are made of little configurations. But when you analyze in just the right way, you suddenly realize that all configurations, of whatever size or kind, are made of sensations, which are not configurations of any kind. And you realize that if it weren't for the presence of those sensations, there couldn't be any configuration to see: a field consisting of a single sensation, such as white, can't lead to any sense of configuration. There's a relationship between these levels of perception. That gives us a hint about building models of perception, a hint about how the brain's perceptual system is constructed.

Sometimes you will identify what seems to be a higher level of perception, some characteristic common to all perceptions, unconnected to lower levels you have previously seen. Then you can use this kind of analysis to try to fill in the gap. What is this new perception made of, besides smaller perceptions of the same kind? When the gap is large, the missing steps are obvious. You can, for example, look at spatial relationships such as "on" -- something being "on" something else. You can see the on-ness clearly, it's right in front of you. But what is it made of? If you said "sensations" you would clearly be making too large a jump, because on-ness involves objects, things, configurations. Some kind of object is "on" some other kind of object. If it weren't for the impressions of distinct objects, there couldn't be any sense of the relationship between them. But is that step small enough? I've had to put two levels between relationships and configurations: transitions (which can be zero), and events (which can be as simple as mere duration). Seeing something "on" something else involves more than a brief contact; there must be duration.

Perhaps someone else could find smaller steps still, or would characterize the intervening steps differently. There's still a lot of room for improving the definitions of the phenomena we're hoping ultimately to model.

I'm not talking here about the models themselves. I'm talking about the attitude you take toward your own experiences when you're trying to notice phenomena that need modeling. If you were a physicist you wouldn't be taking this attitude. You'd treat the world of perception in the normal unanalytical way as if it lay outside yourself where everyone could see it, and you'd search for laws relating changes of one kind of perception to other kinds of perceptions. You would call these "natural laws" or "behavioral laws" and assume you were discovering truths about an objective universe.

As a CT psychologist, however, you have a different objective: to grasp the natural world as a manifestation of human perception (your own), and to ferret out of it some regularities that tell us about perception rather than about the world perceived. If you stumbled onto this attitude accidentally, without understanding what you were doing, you might well find yourself in a state with a clinical name: dissociation. I don't recommend this attitude as one suitable for ordinary living. It's difficult and uncomfortable, and it tends to strip the meaning from experience (until you get past a certain point, after which you realize that meaning, too, is perception, and let it back in). If you're afraid that understanding your girl friend as a set of intensities, sensations, configurations, transitions, events, relationships, categories, sequences, programs, principles, and system concepts in your brain might strain your feeling toward her (and hers toward you), don't do this with your girl friend. Do it with somebody else's, or a laboratory rat. It doesn't matter who or what you do it to, because you're really talking about your own perceptions. This is a private experience valid only in one person's world. It can become public only to the extent that different people independently arrive at the same analysis. I've always hoped for that, but only a very few people, to my knowledge, have tried this for themselves. Most people just memorize my definitions, which unfortunately are in words. It's easier to push words around than to shut up and examine direct experience.

You'll hear objections to this process alluding to introspectionism, which failed to get anywhere a long time ago. But introspectionism didn't fail because it looked at the kinds of things I'm talking about here. It failed because it confused the subjective with the objective (and so did its critics). The world that I'm speaking of examining here would be called, by most conventional scientists, the objective world, not the subjective one. I'm not recommnending shifting attention off the objective world and plunging into the dim and uncertain world of inner phenomena -- or what we imagine to be inner phenomena. I'm recommending a change of attitude toward the world we normally consider to be the objective one, which includes the world outside us and our bodies as we experience them. I'm saying that you will learn something if you look on this world as directly experienced evidence about the nature of your own perceptual system, and only in a conjectural way about the world that is actually outside you.

Instead of treating relationships like on, beside, after, with, and into as properties of the external world, look on them as perceptions constructed on a base of lower-level perceptions. Instead of seeing categories as made of things that are inherently alike, think of categories as ways of perceiving that MAKE things appear to be alike -- things that are actually, at lower levels of perception, different. Instead of seeing sequential ordering as a fact of nature, see it as a way of putting ordering into an otherwise continuous of miscellaneous flow. In short, take nothing about experience for granted, as if some aspects of experience were really outside and others were inner interpretations. Put the whole thing inside, and see what you come up with when you understand that it's all perception. All of it.

Final notes:

In HPCT diagrams, we show signals coming out of perceptual functions and going into higher-level ones (as well as the local comparator, if the signal is under control). I think of these lines as representing single neural signals that vary in only one dimension: how much. This can be confusing, because we don't experience single signals under normal circumstances (when we do they cease to be meaningful). Instead we experience all the signals within the scope of awareness, at every level in the state we call conscious. To understand what the single-signal concept means, you have to break this world of simultaneous perceptions into its components, the individual and independent dimensions in which the totality of perception can vary. You have truly identified one isolated perception when it can vary only in the degree to which it's present, which we experience as its state. If the perception varies without in the slightest changing its identity, you have probably noticed a single signal.

This can be important when you talk about control. We talk loosely about controlling "a dog," for example. But that way of talking is really lumping many independently variable aspects of the dog together. You don't control its species, or its eye color, or the length of its tail. You don't even control its behavior. If it's behavior you're controlling, you always control SOME PARTICULAR VARIABLE ASPECT OF THE DOG'S BEHAVIOR. You may control the radius within which it can move, by putting it on a chain. You may control its speed of walking by saying "stay" or "follow," and its path by saying "heel." Whatever you control, it must come down to a single variable or small sets of variables independently controlled. If you're controlling in more than one dimension, you must sense more than one variable, and have a control system operating independently for each one. That's because independent dimensions can be independently disturbed; you need independent control systems so that a disturbance in one dimension can be corrected without necessarily causing an error in another dimension.

None of this answers your question as to how perceptual signals in a diagram depend on perceptual signals lower in the diagram. The only general answer I can give is that some computation lies between them. The input data consists of lower-level perceptions; the output data, the higher-level perceptual signal, represents the value of the function being computed over and over or continuously. At each level, I presume (judging from the way the context changes every time you consider a higher level), a new type of computation is involved, not simply a repetition of the kind of computation at the lower level. The process of deriving categories from sets of relationships can't be carried out by the same kind of computation that derives relationships from sets of events or lower perceptions. There is no one kind of computation that could serve at all levels.

But as I say, I am, we all are, a very long way from grasping what these kinds of computations are. Every time people come up with a new computer program for recognizing objects, they try to establish this new computation as the blueprint for the whole perceptual system. This is a waste of time. The blueprint changes with every level. Weighted algebraic summation is simply not going to suffice to model our capacity to recognize and execute a program described in words: a rule. Even though such networks are purported to recognize categories, I think that the categoryness is read into the results by a human observer. I don't think that any category-recognizing back-propagation model will actually create what human beings experience as categories -- for example, the category "wife." Of the eleven levels of perception in my model, I think we know how model two of them, the first two. All the rest of our modeling presents to us what a human being might recognize as a higher-level perception, but which the circuit or program itself does not recognize -- or control.

In that I could be wrong, of course, because I speak the truth when I say I don't know how the higher levels of perception work. That means I don't know how they don't work, too. I'm just expressing a hunch.

It's late and I've posted this so I could get to sleep (some ideas just have to leak out through the fingers before they'll let you alone). I'll get to comments on other interesting mail tomorrow.

Best,   Bill P.

Date: Tue Mar 24, 1992 3:44 am PST
Subject: Re: "completely" closed loop = ?

[From Chris Malcolm]

Rick Marken (920323 8:30) writes:

>I just can't understand how any "behavior" that consistently
>achieves some result (escape) could be completely open loop (unless it were
>always performed in the same environment, from the same orientation, with the
>same motor charateristics, etc.). The achievment of consistent results
>in a variable environment is control, by definition.

This is a common misconception. It is possible for behaviour to be adaptive in the sense of changing appropriately in response to the environment without even involving any sensors, let alone controlled variables. There are two ways it can be done. The first is to use unstable motor behaviour where the environment selects the appropriate behaviour. The second is where the appropriate behaviour concerns part of the environment rather than the creature, e.g., putting something into a hole. This can be done by using fixed predetermined behaviour on the part of the creature to determine only part of the behaviour of the item in question, and letting the local environment also affect it, with the net result that fixed behaviour on the part of the creature results in goal-seeking behaviour on the part of the item in question (the thing which is being put in the hole, for example).

These two cases seem to many people to be so counter-intuitively absurd that they cannot believe it possible, but as the history of science has so often shown, failure of our human imaginations is not a reliable guide to impossibility in our universe!

In order to convince students of this I have built a couple of simple robot systems which demonstrate these two cases. Look Ma, adaptive behaviour with no sensors!

After a suitable pause for speculative ingenuity I will describe these systems.

Only two ways? Well, I have only been able to think of two ways. I hope some of you will be able to think of some more!

Date: Tue Mar 24, 1992 5:12 am PST
Subject: RE: Bill on Levels

From Pat Williams (920324)

I really liked your explanation of "Levels of Perception," Bill. Hearing how you arrived at the levels you have found, rather than just names and descriptions of the levels makes it much clearer to me. It is really fascinating to me to think about how perceptions can be combined to form higher level perceptions. I have a hunch that computer programmers (at least good ones like you) may be better at this kind of thinking than most people, since they have to break things down to minute details to make anything work. Just determining the simplest perception like edge recognition is amazingly complicated. I'm currently working on an automatic curve tracer for PictureThis. Determining the edges, corners, and intersections of curves when you only have local pixels to work with is incredibly difficult. It seems fairly trivial until you try it and find all the exceptions. And of course that is no where near as complicated as what you are trying to figure out.

Best wishes,   Pat

Date: Tue Mar 24, 1992 10:14 am PST
Subject: address correction

pP[From Bill Powers (920323.0900)]
The address is, of course, powers_w%flc@vaxf.colorado.edu
The first name is william, not denison.
Sorry.

Bill P.

Date: Tue Mar 24, 1992 10:26 am PST
Subject: The nature of control
From: Randy Beer

Based upon Bill Powers' and Rick Marken's replies to my previous posts, it seems to me that we are simply not communicating. Before I turn to some of the specific points, let me ask a series of general questions in another attempt to understand your position.

1) Can a autonomous dynamical system (one w/o any inputs) ever be a control system? I suspect that your answer would be no (if I'm wrong, then please tell me). However, I would say that it could. Consider a washing machine. As far as I know, the control circuitry of a washing machine typically does not use any sensors, yet it causes the motors and valves to operate in such a way that clothes are cleaned. For the same reason, a purely central pattern generator can also be a control system.

```
-------------            -------------
|           |            |           |
|Controller | =========> |   Plant   | =========>
|           |            |           |
-------------            -------------
```

2) Can a nonautonomous dynamical system (one with inputs) that does not employ any feedback ever be a control system? Again, I would say yes. Consider the washing machine. By pushing different buttons or turning different knobs, I can make it wash my clothes in different ways. Another example of this might be the moth that folds up its wings when it detects a bat's sonar signal. To me, this like a nonautonomous dynamical system that switches modes when triggered by a particular input transition. I personally find it very strange to talk about this as a negative feedback system in which the moth's perception of the bat's sound is controlled. Moth's can hear the bat's sounds for great distances. Once it begins to fall, this perception doesn't suddenly go away. Also, if the bat continues to pursue the moth, it does not take any further evasive action in an attempt to minimize the deviation of its BAT-ATTACK signal from the desired reference level of zero. It's simply been "wired up" by evolution to stop flying when it hears a sound of a given frequency.

```
           -------------            -------------
           |           |            |           |
=========> |Controller | =========> |   Plant   | =========>
           |           |            |           |
           -------------            -------------
```

3) Can a nonautonomous dynamical system with noncontrolled feedback ever by a control system? This one may be a little subtle, so let me explain what I mean. By noncontrolled feedback, what I mean is that, though the system receives feedback, it is not the purpose of the controller to control that feedback (i.e. disturbances in the feedback are not compensated for). Rather, the purpose of the system is to control other outputs which are not directly fed back into the controller. Again, I would say such systems are control systems. Note that I will grant you that the presence of any feedback complicates the analysis of the complete system. But this isn't anything new to neuroscience, since nervous systems have a great deal of INTERNAL feedback, let alone the EXTERNAL feedback through the world.

```
           -------------            -------------
=========> |           |            |           | =========>
           |Controller | =========> |   Plant   |
      ||==>|           |            |           | ========||
      ||   -------------            -------------         ||
      ||                                                  ||
      ||==================================================||
```

4) Finally, we have nonautonomous dynamical systems with controlled feedback, which I belive all of us would agree are control systems.

What I am trying to argue is that the concept of control is more general than negative feedback control. It includes other kinds of feedback (even positive feedback might be involved, such as in the feeding controllers of the artificial insect and the marine mollusc _Aplysia_. Positive feedback doesn't result in divergent dynamics because neurons saturate). My notion of a control system is any dynamics that can cause some plant to respond in a desired way. A room full of motors, pipes, etc. won't wash your clothes. But if you hook them up in the right way and add the appropriate CONTROL (even w/o any feedback whatsoever!), then you get clean clothes.

I have talked to some of the people I collaborate with in the systems engineering department (our collaboration involves the evolution of control systems using GAs and dynamical neural networks) and they agreed that they would consider all of the above scenarios to be control systems. So I will be very interested to hear your answers to the above questions.

Both of you seem to place a great deal of significance in the claim that the only things that matter are the controlled variables, everything else is just in the eye of the observer, an "irrelevant side effect". I simply cannot understand this position. In the course of controlling for, say, its velocity, a moth may rip off its wing. This may be irrelevant to you, but it is certainly not irrelevant to the animal. THINGS THAT ARE NOT EXPLICITLY CONTROLLED (in a closed loop, negative feedback way) MAY STILL BE OF THE UTMOST IMPORTANCE TO THE SURVIVAL OF AN ANIMAL AND THEREFORE SELECTED FOR (or against, depending upon whether they increase or decrease the survivability of the animal).

Now on to some of the specific issues that your replies raised:

First of all, I would like to assure both Bill and Rick that I am not arguing for the superiority of stimulus-response explanations over control theory explanations, though I believe that S-R explanations are sometimes quite appropriate. I would also like to state that I do not believe that feedback is inherently too slow to be of any importance in biological systems. Feedback is quite obviously crucial in many biological systems, but it is not universally necessary.

In claiming that feedback is in continuous operation, I think that you may be forgetting a very crucial fact about nerve cells. Only action potentials propagate any significant distance. Any subthreshold variations in the membrane potential of one cell do not, in general, affect other cells (there are important exceptions to this, but these exceptions are not related to my response below). So I hope that you would agree that, if some sensory organ does not even fire an action potential until AFTER the event it's supposed to be controlling is over, then feedback can play no role. By the way, I would also like to point out that action potential propagation is usually one of the smaller components of delay in the nervous system. The speed of transmission in chemical synapses, membrane time constants (which affect the subthreshold responses before an action potential is fired), and the responses of the sensory structures themselves can all make significantly larger contributions.

Both Rick and Bill questioned my claims regarding the speed of feedback in cockroach escape and walking. With the above observations, let us turn from armchair speculations about how biology ought to work to the biological data. To keep things simple, I will just describe Zill's claim that, while sensory feedback is important at slow speeds of walking in the cockroach, it plays no significant role at high speeds of walking.

Zill was studying a variety of sensory organs in the cockroach leg.

The roles of these various sensory organs play during cockroach walking are known. For example, there are groups of organs known as campaniform sensillae (CS) which are sensitive to stress in the cuticle along a number of different directions. Different groups of these organs are known to be involved in initiating and terminating the stance phase through their connections to a major extensor muscle in the leg. During slow walking, the bursts in one group of CS (the proximal CS) immediately precede the burst of activity in the extensor muscle, while the bursts of another group (the distal CS) immediately precede the termination. This phasing is not accidental; direct stimulation of the proximal and distal CS groups have been shown to produce reflex affects consistent with the functional role described above.

During fast walking, however, the phase of bursts in the CS shift significantly relative to the extensor bursts, so that the proximal CS burst about 20 msec AFTER the beginning of the extensor burst that they normally initiate (and likewise for the distal CS and the end of the extensor burst). This fact would seem to make it difficult to argue that these sensors are playing any role in influencing events that are over before their influence even begins.

There is also some more indirect behavioral evidence that sensory feedback does not play any significant role in fast walking. At slow speeds of walking, the individual leg movements are quite variable, while at high speeds, they are very stereotyped. Leg amputation experiments are also relevant. At slow speeds of walking an amputee changes its normal leg movements so that that its center of mass is always supported. However, at high speeds of walking, an amputee reverts to the normal tripod pattern of leg movements, causing instability and frequent slipping and falling.

This response is already getting too long, so let me leave most of your Evolution/Reorganization comments for another time. Let me respond to just one point. I understand that a biased random walk is not EQUIVALENT to gradient descent, for exactly the reasons that you state. My only point was that IF you used gradient information to bias your random walk, THEN it would be identical to gradient descent. In fact, your notion of reorganization sounds very much like a search technique called random search, with a variable mutation rate. We have been experimenting with a variety of search techniques in addition to GAs, including gradient descent methods, simulated annealing, and random search.

In summary, I am not arguing that your notions of closed-loop, negative feedback control and reorganization are wrong. Quite the contrary, I think that negative feedback and plasticity play very important roles in animal behavior. But I do not believe that they even come close to exhausting the available mechanisms. In particular, I think that you underestimate the role of autonomous, feedforward, and noncontrolled negative feedback dynamics in control. I also think that you underestimate the role of evolution (and development, another whole process that intervenes between genotype and phenotype) in the design of nervous systems.

Regards,   Randy

Date: Tue Mar 24, 1992 11:01 am PST
Subject: Malcolm's automata

From Greg Williams (920324)

I think Chris is providing his no-sensor "adaptive" robots with environments which are disturbance-free with regard to their "adaptations," and therefore he doesn't really counter Rick's argument, which is: IF disturbances are acting to prevent an outcome, THEN control via perceptual feedback is the only way to reliably achieve the outcome (although the outcome might be achieved SOMETIMES -- "by accident" -- without such control).

One might call what Chris is talking about "environmental guidance." Tracks, for example, keep trains where we want them. Note the "we"! I suppose that Chris has designed environments for his robots to keep the robots doing what he wants. And in doing the designing, he eliminated the possibility of disturbances acting to prevent the outcome he wants. (Give me a screwdriver and two minutes with one of your environments, Chris, and I'll show you what a REAL disturbance looks like!) So it really boils down to "creator's guidance," as with clockwork automata of other sorts. They ARE adaptive -- to their creators' desires! Those desires were achieved via control of perceptions by the creators.

Greg

Date: Tue Mar 24, 1992 11:14 am PST
Subject: Purpose, Failure of PCT

[From Rick Marken (920324 8:20)]

First -- a note to Chuck Tucker. I read your "study guide" on sociological views of crime that you sent to me by snail mail. Thanks so much. It was really excellent. Not only well written -- but I could even warm up to the conclusion. I'm surprised to find that Durkheim was such a marvelously innovative thinker. Really nice paper Chuck.

Jeff Dooley (920323.1100)

Welcome to CSGNet.

Your comments on "purpose" and "open loop" behavior were right on target. Kuhn would have a great time with this one all right.

You also say:

>I'd like to suggest the idea, blending concepts of control
>theory purposiveness with a thesis of biologist Stuart
>Kauffman's: that species (complex, adaptive systems) evolve
>(drive themselves) to the edge of chaos and maintain
>themselves there.

I think I need some clarification of the "edge of chaos" idea -- maybe you could make it a bit more concrete. It seems to have a family resemblance to the idea that evolutionary reorganization

(which seems a bit chaotic) is the result of an increased rate
of random mutation (over generations, of course) resulting from
chronic "error" at the genetic level.

Chris Malcolm, in response to the following comment by me:

>>I just can't understand how any "behavior" that consistently
>>achieves some result (escape) could be completely open loop (unless it were
>>always performed in the same environment, from the same orientation, with the
>>same motor charateristics, etc.). The achievment of consistent results
>>in a variable environment is control, by definition.

says:

>This is a common misconception. It is possible for behaviour to be
>adaptive in the sense of changing appropriately in response to the
>environment without even involving any sensors, let alone controlled
>variables. There are two ways it can be done. The first is to use
>unstable motor behaviour where the environment selects the appropriate
>behaviour.

That's one smart environment you got out there, Chris.

>The second is where the appropriate behaviour concerns part
>of the environment rather than the creature, e.g., putting something
>into a hole. This can be done by using fixed predetermined behaviour on
>the part of the creature to determine only part of the behaviour of the
>item in question, and letting the local environment also affect it, with
>the net result that fixed behaviour on the part of the creature results
>in goal-seeking behaviour on the part of the item in question (the thing
>which is being put in the hole, for example).

I don't know whether to laugh or cry. First I read this; then I go to
the mail box and find that the symposium for APS that I proposed was
rejected (Bill, Joel and Tom -- you're welcome to come to San Diego and
party at my condo down their anyway; you can leave your senses at home,
of course, because all we'll really be doing down there is behaving --
adaptively, I hope, but my experience has been that the San Diego
environment gives great adaptive behavior -- if you consider playing golf
and swimming to be adaptive).

Imagine my disappointment when I learned that my concept of behavior as
the control of perception is a common misconception; and on the same day
that I find out that the major national association of scientific
psychologists has no interest in a symposium on the implications of
purpose for the study of behavior. Boy, there's 12 years down the tubes.
I'm sorry all you folk's on CSGNet had to waste all this time and
bandwidth on arguments over a common misconception. But at least Chris
left the door open to the possibilty that SOME "adaptive" behavior
involves sensory input. So there are controlled variables -- they just
don't really need all this control theory stuff to understand them.

Chris goes on to say:

>These two cases seem to many people to be so counter-intuitively absurd
>that they cannot believe it possible, but as the history of science has
>so often shown, failure of our human imaginations is not a reliable guide

>to impossibility in our universe!

Well, it was my impression that many people feel these two cases to be so obvious that they are the basis of all studies in the life sciences. Guess I run into the wrong people.

>In order to convince students of this I have built a couple of simple
>robot systems which demonstrate these two cases. Look Ma, adaptive
>behaviour with no sensors!

>After a suitable pause for speculative ingenuity I will describe these
>systems.

Please -- don't pause for too long. Seal my doom -- drive the stake into the heart of PCT. Describe for us, in detail, these examples of adaptive behavior (by which I presume you mean control since that was what I was referring to in the original statement) that are open loop.

But just one teensy-weensy question before you do this -- so you can be sure that your aim is true and you go straight to the heart: in your "putting item into the hole" robot -- does it work no matter what the size and shape of the item (even if it's wider than the hole)? Does it work on a windy day? If the robot moves or is pushed (with continuously varying force)? If you block the hole?

If "item in hole" is the reference state of a controlled variable then that result should be achieved in the context of any disturbance that could change that result -- to "item next to hole", for example. Does the robot do that? Does it get the item into the hole in spite of reasonable disturbances? If so, we control theorists have sure made a BIG mistake. You will forgive me for thinking that it can't do it -- but if you demonstrate to me that you have a robot that can control a variable without sensing it, then I promise to be a good scientist and admit that I was as "mis-conceptioned" as I could be -- and admit that SR is an important model of adaptive behavior.

One last little point -- I am assuming that you mean the same thing by "adaptive behavior" as I mean by control. If not -- if, for example, all you mean is that you can design a system that generates interesting results (like putting an item into a hole) then we are not talking about the same thing at all (as you should be able to tell by my original statement). I am not claiming the you need to control sensory input in order to produce "interesting" or "useful" results -- you can produce many useful results by accident. I am saying that systems cannot produce CONSISTENT RESULTS -- whether they are interesting or useful or whatever from YOUR point of view -- in normal environments -- ie -- ones which also have variable effects on these results -- unless the systems can SENSE these results. If this is a misconception -- common or not -- I would apreciate being straightened out on it as soon as possible; certainly before I waste anymore trees writing papers about it.

Regards   The loose canon

Date: Tue Mar 24, 1992 12:00 pm PST
Subject: Reorganizing Escape

[from Gary Cziko 920324.1320]   Bill Powers (920322.1700)

>Consider what Beer said about waxing the cerci of the cockroach. This
>throws off the detection of the direction of winds, so the cockroach's
>escape response goes in a wrong direction. That's what you'd expect if the
>escape response were wired in, being changeable only from one generation to
>the next. But Beer then said that if the threatening puff of air is
>repeated again and again, something changes inside the cockroach, the very
>same live cockroach within its own lifespan, such that the escape response
>becomes appropriate again. Obviously, evolution had nothing to do with this
>change, because cockroaches have to die before evolution comes into play.
>The internal machinery of the cockroach somehow detected the
>inappropriateness of the escape response, and altered itself to make it
>appropriate again.
>
>What criterion could have directed this sort of change? The only possible
>one is that the escape response was not having an opposing effect on the
>"bad" stimuli that gave rise to it, over many trials. The change of
>organization stopped only when the escape response was in the direction
>that once again took it downwind, toward smaller velocity fields, and thus
>led to a lessening of the stimulation.

When I saw this I said to myself "Why didn't I think of that" (Rick Marken just claims that he didn't even SEE Beer's description of the reorganization of the cockroach's!). And yet, some people (like Greg Williams) keep trying to push the idea that this is a "ballistic" response to puffs of air.

How can the cockroach change his behavior adaptively in this situation if it doesn't care where its going (e.g., make sure ground speed is faster than air speed, that means I going downwind, away from the air disturber)? Please, somebody explain to me how this can happen if the cockroach's escape behavior is purely ballistic.--Gary

Date: Tue Mar 24, 1992 1:17 pm PST
Subject: Control vcs. noncontrol

[From Bill Powers (920324.1100)]

Chris Malcolm (920323.0830) --

You're not talking about control, Chris, but simple cause and effect. The only control involved in the scenarios at which you hint is your own -- you adjust the environment so that what you want to happen happens.

I can think of many examples of this sort of HUMAN control. Say you want to pour gasoline into a gas tank out of a container that sloshes and blurps and sends gasoline in all directions. How do you build a device that will make sure the fuel goes through the little filler hole, no what its direction of egress from the container? Easy: buy a funnel.

The funnel, however, isn't a control system. If something pushes the small end of the funnel away from the filler hole, the funnel won't push back so as to keep the gasoline going through the hole. If it were a control system, it would.

Real control systems can achieve consistent ends even in real environments where disturbances can affect consequences of the system's output, change the characteristics of the environment, and cause changes in the system's own effector calibrations. If you don't take unpredictable independent disturbances realistically into account, you're not designing control systems, but 19th-century automata. I have no doubt that such automata can be made to do very clever things very precisely -- as long as they are built so natural disturbances and changes in their own characteristics are prevented. This is how all machines were built prior to control systems. Up to a point, it works. But this is not how organisms work in real environments, and it's not how any control system to which the name properly applies works.

Pat Williams (920324) -- Such words of appreciation, especially from you, are high praise indeed. Thank you.

Randy Beer (920324) --
>1) Can a autonomous dynamical system (one w/o any inputs) ever be a
>control system? I suspect that your answer would be no (if I'm wrong,
>then please tell me). However, I would say that it could.

Then we mean something different by "control system." You aren't the only one who differs with me on this, but perhaps I can explain my usage of this term in a way that will make my claim more palatable.

A control system can do something that no conventional device can do: produce a consistent outcome under conditions where (1) varying actions are required to produce a consistent outcome, and (b) its effector calibrations are subject to unpredictable drifts. This is accomplished through sensing the state of the outcome directly, comparing the result of sensing with some reference criterion, and using the result of the comparison to adjust the drive to the effectors. This can result in very precise control of the outcome, even if the effector sensitivity to the driving signal varies over a factor of ten and even if there are external influences that can affect the outcome just as much as the effector can. What I mean by a control system is a system that can work this way.
Let's take your Fig. 1, with an addition by me:

```
-------------            -------------
|           |            |           |
|Controller | =========> |   Plant   | =========> variable
|           |            |           |              ^^
-------------            -------------              ||
                                                    ||
                      Unpredictable disturbance ====
```

A controller with or without an input (but without feedback sensing of the variable) can't maintain the variable at any particular value as long as the disturbance is present. Neither can it do this if the plant's response to the signal from the controller varies in amount. In order for a system like this to achieve stability of the variable against direct disturbance, the "plant" part must be built so massively and be so rigidly coupled to the variable that the disturbance simply has no significant effect.

There's another factor that lies quietly in the background, unspoken.

That is the human manipulator who adjusts this controller and plant so that (without disturbances) the variable comes to the correct state. The human manipulator can't do this without sensing the state of the variable and comparing what is observed with a desired state.

Adjustments of the controller and plant are made until the observed state of the variable matches the desired state. If the properties of the controller or plant drift so the variable departs from its desired state, the human controller will see a difference between the actual and desired state, and depending on its size and direction, make appropriate adjustments to plant and controller. So the loop is always closed if real control exists.

>2) Can a nonautonomous dynamical system (one with inputs) that does
>not employ any feedback ever be a control system? Again, I would say
>yes.

And for the same reasons as before, I would say no. The system you draw can't counteract disturbances or compensate for changes in its output properties. The washing machine will happily go through its cycle even if the waterline is clogged (although some washing machines use feedback control, and won't proceed until the proper weight of water and clothes is detected) or the timer sticks on the fill cycle. An open-loop system can produce a consistent result only if there are no disturbances acting directly on the result, and only if its output characteristics remain exactly the same.

>I personally find it very strange to talk about this as a negative
>feedback system in which the moth's perception of the bat's sound is
>controlled. Moth's can hear the bat's sounds for great distances.
>Once it begins to fall, this perception doesn't suddenly go away.

I will have to learn more about how these moths actually behave at different levels of bat sound. If a cockroach can steer by small differences of odor in an inverse-square odor field (its behavior strongly affecting those differences), why can't a moth behave so as to vary the intensity of a bat-sound in an inverse-square sound field? But I don't want to make a big case of this behavior -- perhaps it works just as you say it does. I am really more interested in how the moth works the other 99.9% of the time, when it isn't showing any dramatic "responses," but is probably controlling a hundred variables continuously. Why ignore the huge number of control behaviors that are going on every moment that the moth is active in favor of a few unusual seemingly open-loop responses?

>3) Can a nonautonomous dynamical system with noncontrolled feedback
>ever be a control system?

No, once again for the same reason. Consider your diagram, again with the same addition:

```
                                                         disturbance
                                                              ||
           -------------            -------------             ||
=========> |           |            |           | =========>variable
           |Controller | =========> |   Plant   |
      ||==>|           |            |           | ========||
      ||   -------------            -------------         ||
      ||                                                  ||
```

||==||

I'll agree that with the feedback loop coming directly from the plant, you now have a system that can be immune to drifts in the characteristics of the plant and controller, so the output (lower one) becomes a reliable function of the independent input. In fact, the sensed state of the plant's output is now compared with the input and the drive to the controller and plant is automatically varied to keep the difference near zero. The output of the plant, as sensed, is truly under control now.

But the variable affected by that output is not under control. Disturbances that affect the variable will simply add algebraically to the plant's output and the variable will assume the resultant state. If there are changes in the link between the plant's output and the variable, the variable will again change. There will be no action to bring the variable back to the undisturbed state. So this kind of system can work only in an unreal or protected environment in which such disturbances can't happen -- or else if the variable is so tightly coupled to the plant's output that it can't vary from the state fixed by that output even when disturbances are present.

Control is not required when a variable affected by a system's output is never subject to effective disturbance, and when the output effectors retain perfect calibration. Systems that can work ONLY under such conditions are not control systems, by my definition.

>What I am trying to argue is that the concept of control is more
>general than negative feedback control.

I know that the term "control" is used in many circumstances where I would not use it. I'm trying to promote a more technical usage of this term, and through this usage a wider understanding of the tremendous differences between systems with and without negative feedback control. This isn't terribly important in engineering, where high precision and massive construction can achieve predictable results without feedback, and where the main thing is to get the job done. But it is important in modeling organisms, because organisms have effectors with very sloppy properties, and the external effects they have on the environment are subject to all sorts of disturbances that can neither be sensed at their sources nor predicted. Without the concept of control that I espouse, it is simply impossible to explain how organisms manage to produce repeatable consequences in the presence of variable disturbances.

>My notion of a control system is any dynamics that can cause some plant
>to respond in a desired way.

And if it doesn't quite respond in the "desired" way, what do you do? You adjust the dynamics. The system will then continue to produce the desired response until disturbances change or the dynamics of the system drift. Then you have to adjust it again. To get the system to produce a consistent response, you have to attend continuously to it, and substitute your own capacity for negative feedback control for the capacity you haven't put into the system. That's why, in photographs of Nineteenth Century machinery, you will often spy the operating engineer lurking in the background, wrench in hand. Control systems don't need such babying.

In computer simulations you don't have these problems, because simulations don't drift and you don't program in random disturbances of the outcome (unless you're doing control-system models). So the simulations seem to work just fine. They would not work in a real environment unless they used negative feedback.

>I have talked to some of the people I collaborate with in the systems
>engineering department (our collaboration involves the evolution of
>control systems using GAs and dynamical neural networks) and they
>agreed that they would consider all of the above scenarios to be
>control systems.

One of the great disillusionments of my life was the discovery that even real control engineers don't have a very good grasp of the differences between control systems and other kinds. To most of them, for example, it comes as a surprise to realize that control systems control their own feedback signals, not their outputs. This isn't what they were taught, but a moment's thought will show that it's true.

>I simply cannot understand this position. In the course of controlling
>for, say, its velocity, a moth may rip off its wing. This may be
>irrelevant to you, but it is certainly not irrelevant to the animal.

You are citing a relevant side effect to refute an observation about irrelevant side-effects. In your example, an effect of one control system's action disturbs something else of importance to the organism. But the importance of the moth's losing the wing is not the same to the moth as it is to you. To you, the primary effect is that the moth can no longer fly and will probably starve. To the moth, "flying" and "starving" are not variables which which it can be concerned. Only the effects on its internal state of not flying and not eating are important to it.

In a less trivial example, a cockroach's path to the food patch may result in the movement of its image across your retina. This is an irrelevant side effect, because it has no effect on the internal state of the cockroach. How YOU see the cockroach moving is irrelevant to the cockroach.

> THINGS THAT ARE NOT EXPLICITLY CONTROLLED (in a closed loop, negative
>feedback way) MAY STILL BE OF THE UTMOST IMPORTANCE TO THE SURVIVAL OF
>AN ANIMAL AND THEREFORE SELECTED FOR (or against, depending upon
>whether they increase or decrease the survivability of the animal).

I think you are taking too limited a view of what constitutes negative feedback. Also, you are not thinking in terms of a hierarchy of control systems, but only at one level. In a human being, arm position is under direct negative feedback control. By varying the reference signal for arm position, however, a higher-level system can cause the arm to reach out and touch a target -- controlling the distance between fingertip and target, as seen. And a higher system still could vary the x and y reference conditions for the relationship between finger and target to make the finger trace a circle around the target, or a square, or any other figure -- again, under feedback control. I can't prove it, but I think that even in natural selection there are feedback control processes involved, in which the organism's actions control the effects of selection pressures on the organism, thus effectively controlling the course of natural selection. You indicate some degree of agreement with this in your post.

>Feedback is quite obviously crucial in many biological systems, but it
>is not universally necessary.

Good. I can agree that not every aspect of behavioral mechanisms involves feedback control. For example, the response of a muscle to a driving signal does not entail, as far as I know, any feedback that modifies the driving signal where it enters the muscle. It isn't necessary even that the larger stretch or tendon reflexes exist, if the position of an animal's limb reliably depends on the signal entering the motor neuron: the damped mass-spring properties of the leg might suffice. However, if it turns out that applying a force to the leg results in an opposing change in the muscle tensions, then feedback control is clearly present. Cruse mentions in one article that this is true of the cockroach. When a clear feedback control effect is observed, I don't think the use of control theory is optional any more, and S-R theory is ruled out.

>I hope that you would agree that, if some sensory organ does not even
>fire an action potential until AFTER the event it's supposed to be
>controlling is over, then feedback can play no role.

This isn't true. Ask your consultants about sampled control systems and z-transforms. In control-system models of neural systems, the variable of interest is frequency of firing. Also, the physical actions that take place happen on a very slow time-scale relative to the scale on which a single impulse is important (particularly when you consider all parallel pathways that carry the same kind of information around the same feedback loop). While an action is getting under way, neural frequencies can change and be changed by the action: there is no such thing as an "instantaneous" response. All the smoothing that occurs makes neural control systems continuous on the time-scale that matters.

In considering whether feedback has an effect, you have to consider not just a single impulse-event, but recurring events. Feedback does not have to be instantaneous to be effective.

>The speed of transmission in chemical synapses, membrane time constants
>(which affect the subthreshold responses before an action potential is
>fired), and the responses of the sensory structures themselves can all
>make significantly larger contributions.

True. But these are the same averaging and smoothing effects that make frequency, not the single impulse, the measure of choice for neural systems. The delays of which you speak are not transport lags, but integration lags. Integration lags have entirely different effects on closed-loop systems than do pure time-delays.

>During fast walking, however, the phase of bursts in the CS shift
>significantly relative to the extensor bursts, so that the proximal CS
>burst about 20 msec AFTER the beginning of the extensor burst that they
>normally initiate (and likewise for the distal CS and the end of the
>extensor burst). This fact would seem to make it difficult to argue
>that these sensors are playing any role in influencing events that are

>over before their influence even begins.

In my attempts at cockroach simulation, I'm using a central pattern generator similar to yours, with the limit detectors simply triggering reversals when the leg angle is a little less than it would be if the pattern generator alone determined the amplitude. At higher frequencies of walking, presumably the amplitude of leg movement will be less. I suppose that if it's enough less the limit detectors won't fire until the central pattern has already reversed.

This is OK with me. I'm not trying to model the pattern generator as a control system. It's just the output function of a control system (of several of them), without internal feedback of its own. Just like a muscle, which works without local feedback. There's always some level of organization at which you won't find any feedback. When you reach that level, you're looking at components of a control system, not whole control systems. You don't have to prove to me that there are organizational units in organisms that work without feedback control. All the control systems I design have at least three such units: an input function, a comparator, and an output function.

On the other hand, the output of the pattern generator for each leg will provide a reference signal for a leg-position control system, with feedback from position sensors. This is called for by the data, which say that a cockroach dragging a weight increases its muscle forces. So each leg will have a position control system, even though the driving signals are coming from a pattern generator that inside itself has no negative feedback control.

Today, by the way, I got a four-neuron pattern generator for one leg to work so that as a speed signal varies over its range, the speed of movement slows down, stops, and reverses, the appropriate reset signals occurring automatically without any need for gated circuit-switching. The swing phase duration, as per the diagrams in your book, is independent of the stance phase duration in both directions. Only basic Beer neurons are used.

>I think that negative feedback and plasticity play very important roles
>in animal behavior. But I do not believe that they even come close to
>exhausting the available mechanisms.

Nor do I believe that the available properties of control systems have been tapped in the modeling of organisms.

>In particular, I think that you underestimate the role of autonomous,
>feedforward, and noncontrolled negative feedback dynamics in control. I
>also think that you underestimate the role of evolution (and
>development, another whole process that intervenes between genotype and
>phenotype) in the design of nervous systems.

Well, this will all come down to modeling, won't it?

Joe Lubin (920323) -- The arm program went into the mailbox quite a few days ago -- you should receive it any minute. I mailed it the same day you asked for it.

Best to all,   Bill P.

Date: Tue Mar 24, 1992 2:13 pm PST
Subject: The nature of exasperation

[From Rick Marken (920324 13:00)]

Fortunately, I was called away to a big, fun departmental lunch break before I had time to complete my exasperated reply to Beer's very polite post (920324). Now I come back and see that Bill Powers has handled the situation for me with his usual aplomb. Thanks again Bill.

There was a part of the Beer reply I was working on that might be of interest (it was also the only non-exasperated part). Here it is:

Randy says:

> Another example of this might be the moth that folds up its
>wings when it detects a bat's sonar signal. I personally find it strange
>talk about this as a negative feedback system in which the moth's
>perception of the bat's sound is controlled. Moth's can hear the
>bat's sounds for great distances. Once it begins to fall, this
>perception doesn't suddenly go away.

But it does tend to decrease!

> Also, if the bat continues to
>pursue the moth, it does not take any further evasive action in an
>attempt to minimize the deviation of its BAT-ATTACK signal from the
>desired reference level of zero.

Apparently, once it's on the ground there is not much it can do but stay still. Does it curl into a ball on the ground if the sensed sound get's loud enough? If it does -- and this curling has NO effect on SENSED sound level, then you are looking at the pure, SR system. This makes me realize that the exact same input-output function can be feedback control in one situation and SR in another. It is the existence of the simultaneous R-->S FUNCTION that makes the system a control system. If there is no R-->S function, or if it is near zero then for all practical purposes it is an SR relation-ship. So -- if the moth really does curl into a ball in response to sensed sound when it's on the ground just as it does when flying; and if the ground curling has virtually NO EFFECT on the sensory cause of the response, then, IN THAT CASE, the moth's behavior is, I think, SR.

I just can't believe that the moth actually acts this way; if it curls on the ground then maybe the bat could can detect its movement-- or, if the moth actually does respond to sensed sound that way when on the ground, then I would think that the curling would fool or have some other influence on the bat so that the bat tended to wander off (and thus, reduce the sensory input that caused the curling).

But let's assume that it does respond to sound the same on the ground as in the air -- AND that the curling, in that case, has no effect on sensed sound.

What this would shop is that the SAME SYSTEM can change from an SR (open loop) to a closed loop control system just as a result of changes in the physics of its environment. Again, I would imagine this is actually rare. But it demonstrates that the unusual properties of a control system (such as the fact that it controls rather reponds to its sensory input) result for the negative feedback loop in which the system normally exists. If you can somehow break that loop (by removing the influence of output on input) WITH NO CHANGE IN THE INTERNAL STRUCTURE OF THE ORGANISM then the organism becomes an SR device. Actually, I did this in one of my studies where I surreptitiously broke the R--S connection from handle to display in a tracking task. The behavior of the subject became SR -- until they noticed that there was something fishy going on (higher level perception) and then realized that they actually were not in control.

How's that for admitting that SR can exist in living systems. It's just pretty rare (and non-existent as a means of control -- I presume we know what CONTROL means now?).

I'm not sending the above to Randy since he was not in on our SR discussion.

Regards   Rickala

Date: Tue Mar 24, 1992 2:20 pm PST
Subject: Mail for mark olson

Will someone please forward today's mail from me to Mark Olson?
I can't get his address to work.
Thanks --(Gary) --
Bill Powers

Date: Tue Mar 24, 1992 2:38 pm PST
Subject: Non-artifactual automatisms

From Greg Williams (920324-2)

In the spring, young spiders of some species crawl up to relatively high places, like the tops of fence-posts, and raise their bodies up while releasing some silk. They are wafted away by the wind, to land perhaps only a few inches or possibly many miles from the starting point. This "ballooning" disperses the spiders. Presumably, being crowded tends to reduce reproductive success, and so ballooning was favored by natural selection.

Assuming that all of the spiders have essentially identical, non-modifiable (by the spider) aerodynamics, ballooning is analogous to the "adaptive" actions of Malcolm's automata and to the clothes-cleaning of Beer's "feedforward" (or I suppose it could be termed timing-cycle) washing machine.

In my previous post today, I claimed that Malcolm's automata are able to do what he wants them to do RELIABLY because he constructed the automata's environments so that there are no relevant disturbances. I further claim that Beer's timing-cycle washing machine cleans clothes (with respect to a standard for what is "clean" set up by a person, not by the washing machine!) reliably ONLY if the disturbances are not too great (no very heavy stains, not too many clothes per load, no severe electrical brownouts, not much variance in the chemistry of the soap powder used, etc.). In both cases, the requirements to accomplish what is desired are established in advance: Malcolm tinkered with the automata's environments until the robots did what he wanted, and the washing-machine supervisor figured out how long the wash cycle should last for "clean" clothes. It was possible to settle on stable environments or cycle-timing ONLY because the relevant disturbances didn't change unpredictably for each trial of tinkering or washing a load.

The same is true for ballooning of spiders. This "adaptive" "feedforward" action was selected for because the relevant disturbances stayed within certain limits over a (very) long period of time. Gravity did not reverse sign; the wind continued to blow chaotically. If God had built a wind tunnel and subjected individuals in successive populations of ballooning spiders to a constant wind, so that they ended up essentially non-dispersed, the "adaptiveness" of ballooning would have disappeared, and we can predict that ballooning itself would have eventually disappeared, given some lability of the organization of genetic material.

What I'm trying to point toward is a notion that organisms can indeed have "feedforward" "adaptive" actions which are selected for/against, as Randy claims, BUT that such activities are to be expected ONLY when the relevant disturbances remain similar over evolutionary time. If the relevant disturbances are such that they cannot be pre-compensated for, perceptual feedback is needed for reliable accomplishment of the activities.

Greg

Date: Tue Mar 24, 1992 3:10 pm PST
Subject: Cockroach escape reorganization

From Greg Williams (920324-3)

>[from Gary Cziko 920324.1320]

>How can the cockroach change his behavior adaptively in this situation if
>it doesn't care where its going (e.g., make sure ground speed is faster
>than air speed, that means I going downwind, away from the air disturber)?
>Please, somebody explain to me how this can happen if the cockroach's
>escape behavior is purely ballistic.

The precalibrated settings for ballistic escape might be reset on the basis of wind direction/intensity AFTER the escape was completed. No feedback control would be needed during the escape for this to happen.

Greg

Date: Tue Mar 24, 1992 4:32 pm PST
Subject: Re: Malcolm's automata

[FRom Chris Malcolm]

>>From Greg Williams (920324)

>I think Chris is providing his no-sensor "adaptive" robots with environments
>which are disturbance-free with regard to their "adaptations," and therefore

>he doesn't really counter Rick's argument, which is: IF disturbances are
>acting to prevent an outcome,

No I'm not. There is a definition of purpose and disturbance available in which the robots can clearly be seen to be achieving their purpose despite disturbances.

>THEN control via perceptual feedback is the only
>way to reliably achieve the outcome (although the outcome might be achieved
>SOMETIMES -- "by accident" -- without such control).

They achieve their purpose always, without perceptual control, provided the environment is kept within the design limits. Of course their purposes can be frustrated by disturbances they can't cope with, such as setting them on fire or pouring a bucket of sand over them, but the same kind of restrictions apply to perceptual control.

>One might call what Chris is talking about "environmental guidance." Tracks,
>for example, keep trains where we want them. Note the "we"! I suppose that
>Chris has designed environments for his robots to keep the robots doing
>what he wants.

This is true only in the vacuous sense that there are limits of disturbance beyond which the robots can't cope. In practice given an environment and purpose I tried to design a robot to do it. Having succeeded, I then tried varying the environment to find the limits of the adaptability. Then I tried to improve the robot to handle bigger disturbances. And so on.

>And in doing the designing, he eliminated the possibility of
>disturbances acting to prevent the outcome he wants.

As it happens I didn't trim the environment to fit the automaton, I improved the automaton to handle the environment. But this is simply an implementation detail. Given a robot, an environment variable within limits, and a purpose achieved despite these variations, why should the details of the process by which I arrived at the design matter? Suppose I had simply happened by serendipity on a neat trick, and then discovered its scope by experiment, how does that change things? Are you arguing that my automata can only borrow my own purposes? That sounds horribly like the traditional "magic" (intentional, subjective, etc.) view of purpose which I thought perceptual control was supposed to be able to rescue us from?

>(Give me a screwdriver
>and two minutes with one of your environments, Chris, and I'll show you what a
>REAL disturbance looks like!)

Of course you can. Now you show me an insect doing perceptual control and I'll show a boot which can stop it! So what?

>So it really boils down to "creator's guidance,"
>as with clockwork automata of other sorts. They ARE adaptive -- to their
>creators' desires! Those desires were achieved via control of perceptions by
>the creators.

Well, there is an interesting point in here concerning purpose. I am aware of my purpose in posting this, and controlling disturbances which might frustrate it. A dung beetle rolling its dung ball is not aware of its purpose, but it is nevertheless controlling disturbances to the purpose that evolution "designed into" it. So in talking about the purposes of simple mindless creatures (supposing for the sake of argument that dung beetles are mindless) we find we are implicitly talking about the "perceptions" of evolution. Really? What happened to the idea of _objectively_ discovering purpose by _experimentally_ discovering the controlled perceptions?

I applaud that view. My purpose is to argue for its extension, i.e., there are more purposes than those achieved by perceptual control, but they can still be discovered by observation and experiment of a similar kind, i.e., you can discover the purpose of my "senseless" robots by seeing what happens when you change things. You do not need to ask the designer, God, or evolution.

Chris Malcolm

Date: Tue Mar 24, 1992 7:28 pm PST
Subject: Let's get quatitative

[From Rick Marken (920324 18:00)]

I'm at home so here is some header :
I think this discussion of open loops and sr and all this stuff should be making it clear just how exceptionally difficult it has been to get people to start really looking at life from a control system perspective. There is real, serious resistence to the idea that behavior is the control of perception. Watch it happen, philsophers of science -- it shows up as efforts to see input control as irrelevent or a small piece of the behavioral puzzle. The constant efforts to grasp at any straw as evidence of open-loop organization; the moth fall, the slowness of neural impulses, etc etc. Anything that looks in any way like open loop behavior is grasped at with such relish; such a sense of "see there! -- OPEN LOOP! Why the fuss? Why the focus -- when, as Bill mentioned -- you can point to examples of controlled variables by just opening you eyes and watching behavior. People are always repeating results in different conditions -- often wildly different. Watch all those cars consistently staying in their lanes -- how often do you see an accident? People getting where they want to go -- avoiding bumping into unpredictably moving people as they do it. Why all the emphasis on open loop? I think very there is a controlled variable being disturbed. And I think it is the idea that open loop models -- the dominant kind in the life sciences -- MUST BE RIGHT. Otherwise, why the fuss? Or is the tendency to point to purported examples of open loop behavior just a response to the stimulus "control theory"?

Here are some test questions for those of you who think open loop processes are important in some way. Do you think that there are ANY examples of closed loop control? What are they? Do you believe that, when there is closed loop control, that the output variable is controlled by the input variable or that the input is controlled by an implicit or explicit reference variable inside the system? Do you think that most behavior is open loop? Just the simple behavior (reflexes)? Just the big stuff (like solving physics problems)? Why would there be any closed loop control at all if open loop processes seem to work just fine? What is your reason for even reading this newsgroup?

Last point to Chris Malcolm. It sure sounds like you "controllers" are good old fashioned sr devices that produce particular results. You say they are not build to resist the kind of disturbances that Greg Williams wanted to test them with. Could you be a bit more quatitative and give me some idea of the kind of disturbace that will be resisted. Can is resist, for example, force disturbances that are greater than the restoring force of the materials out of which you built the robot? Is the loop gain > 1? Perhaps Bill Powers could suggest the right way to measure how well you robot controls -- if you will describe it in a bit more detail.

Thanks Rick

Date: Tue Mar 24, 1992 9:51 pm PST
Subject: "edge of chaos"

[From: Jeff Dooley 920324.1900]

(Rick Marken 920324.0820)

On Stuart Kauffman's notion of the "edge of chaos."

Rick, first a mention of what I take Kauffman's mission to be. Kauffman appears to argue a case for: (1) a dynamics of spontaneous organization in complex systems--organization arising out of complexity alone (as I read it) and not as a function of some ontogenetic structure or blueprint (like a gene)--and (2) a taxonomy of the "sources" of self-organization in complex systems and an understanding of how such order may enable and/or constrain efficacy in natural selection. (Burian-Richardson) (references below)

Kauffman (Univ of Penn) has elaborated these arguments in his forthcoming, _The Origins of Order: Self-Organization and Selection in Evolution_, Oxford, 1992.

A major thesis of this work, according to reviewers Burian and Richardson, is that the properties of order resulting from the self-organizing process in biological and non-biological systems are independent of selection. Kauffman asks, "Must selection have struggled against vast odds to create order? Or did that order lie to hand for selection's further molding?" ("Sciences of Complexity")

So far as I can gather, it is these properties of order or organization which, in species, drive the ensemble to and keep it at equilibrium (absent fatal perturbations) within a range of adaptability he calls "the edge of chaos." It appears to be precisely Kauffman's point that this edge-of-chaos zone is the one in which (to use a PCT term) the reorganizing system deals most effectively with perturbations and constraints by exhibiting the most resilience under perturbation. So he feels it is no wonder that the organizing function of complex systems seeks equilibrium in this zone. The upshot of this is that selection is not to be held the sole source of order in biology. In fact, Kauffman suggests that in case selection may slightly displace evolutionary vectors of generic order in species, those vectors will still "shine through" as manifest properties, not *because* of selection but *despite* it. It follows from this that ontogenetic variation, for Kauffman, is not the whole (or even the most part) of evolution's story. This appears to be about as monumental a gestalt switch for biology as PCT is for S-R psychology!

Ok, now to your question: what is "the edge of chaos?"

Kauffman offers this thesis: "Complex adaptive systems achieve, in a lawlike way, the edge of chaos ("Sciences of Complexity"). They do this through a process of self-organization whose lawlike progression is a function of the complexity of individuals and ensemble together.

The edge of chaos appears to be a (razor-thin) zone of optimal adaptability within a fitness landscape or vector space. The organizing function of the ensemble seeks equilibrium at this spot. Here is Kauffman's description of the vector space:

". . . complex adaptive entities achieve the edge of chaos because such systems can coordinate the most complex behavior there. Deep in the chaotic regime, alterations in the activity of any element in the system unleashes an avalanche of changes, or "damage", which propagates throughout most of the system. . . The butterfly in Rio changes the weather in Chicago. . .Conversely, deep in the ordered regime, alteration at one point in the system only alters the behavior of a few neighboring elements. Signals cannot propagate widely throughout the system. Thus, control of complex behavior cannot be achieved. Just at the boundary between order and chaos, the most complex behavior can be achieved." (Sciences of Complexity)   Ensembles self-equilibrated at the edge of chaos, therefore, promise the most robust ability for phenotypic change (controlling-behavior!) in the event of deformations (perturbations) in the landscape.

Bill P. has been suggesting recently that some kind of control process may be at work in evolution. As I was reading his posts I was also reading the "Sciences of Complexity" paper and somewhere in there the lightbulbs just started popping.

References:

Stuart Kauffman, "The Sciences of Complexity and 'Origins of Order'," _Philosophy of Science Association 1990_, Volume 2, ed. Fine, Forbes, Wessels, (East Lansing: PSA, 1991).

________, _Origins of Order: Self Organization and Selection in Evolution_, (Oxford: Oxford U Press, 1992 (July))

Richard Burian and Robert Richardson, "Form and Order in
Evolutionary Biology: Stuart Kauffman's Transformation of
Theoretical Biology," _PSA 1990_, Vol. 2.

jeff dooley
dooley@well.sf.ca.us

Date: Tue Mar 24, 1992 10:00 pm PST
Subject: Open loop "control"

[From Bill Powers (920324.2245)]

Chris Malcolm --

OK, Chris, you have us all going now -- time to describe your open loop
control system. What's the principle? What's the implementation?

Best
Bill P.

Date: Tue Mar 24, 1992 11:58 pm PST
Subject: Re: open loop "control"

From Tom Bourbon [920325-0:16]---------
Greg Williams [920324-2]. Ballooning by (Charlotte's?) baby spiders
is an *excellent* example of "feedforward" "adaptive" behavior that is
established and maintained only if there are NO DISTURBANCES to the
relevant consequences of the behavior over evolutionary time. Your
perspective was missing in much of the recent discussion concerning
"open loop control," "S-R control," and similar topics. Certainly
actions can occur that, on the time scale of human perceptions, appear
open-loop, but only if the consequences of those actions FOR THE
BEHAVING ORGANISMS are undisturbed. It does not matter how many
consequences a human observer can see changing if they are not
relevant to the observed organism.
  ****
Chris Malcolm [920324-06:04] You certainly threw fat in the fire
with your remarks! As Bill Powers said in a post a few minutes ago,
tell us more about your robots. Might they be pushing the object
against the floor while executing a routine for movement, a
"strategy" that would result in the object slipping into any
hole that ended up in the proper place?
****
Randall Beer [920324 -- I forget which time] Bill Powers, Rick
Marken and Greg Williams have already commented on your proposed
"control processes," and, as you predicted, they rejected the
idea that the systems, as described, were control systems. I
concur. It appears as though you repeatedly overlooked your
own reference signals (intentions), perceptions and attempts to
act to eliminate error, when you used the systems -- a prime
example being the washing machine, which functioned as an
output device for you, the controller. As Bill Powers remarked,
there need not be, probably ought not be, complete control
loops within the various function boxes in a control loop. But
there will be references,perceptions comparators and the like in

the controller who, or which, uses that output device. Washing machines are parts of the process of "washing clothes" only because people build them and use them to create that end, and the people who use them do so to control their perceptions of the state of cleanliness of clothes. The same considerations apply to our creation of, and use of, many other "control" devices, each of which obviously contains many individual "open loop" components -- thermostatically-equipped air-conditioning systems, cruise-control-equipped automobiles, and the like. And certainly they apply to the many devices we simply manipulate, no matter how complex their clock-work-like innards.

Concerning the now-famous "moth-in-a-bat's-sonar-beam," so that we all might be sparred the embarrassments of overly eager imaginations, could you provide several specific citations of empirical data on the phenomenon of wing-folding and subsequent crashing to earth?

****

Rick Marken [920324--?] So APS shot us down, did they? Might it be because they already had all of the important people set in the program over two months ago, leaving all of the rest of us to hold out hopes in vain? After all, they already have someone to talk on the specific topic of the "rich" soup of "models" available to today'd scientific psychologist. I am certain there will be thorough coverage of PCT in that talk, so they just didn't need anyone else. We would have been redundant, especially since the principles of PCT are so very mainstream, these days!

****

General news of a personal nature ---

Several people from the net have asked whether I am moving. I am. On 1 June, I will make the leap from tenured security into the cauldron of life lived on "soft money." I will join Andy Papanicolaou (who many of you met at the last CSG meeting) in Galveston, Texas. We will be at the Magnetoencephalography Laboratory affiliated with The Transitional Living Community of Galveston and with the Division of Neurosurgery at the University of Texas Medical Branch, Galveston, Texas. I was hired on specifically to work on PCT. I will continue my work on interactins among and between control systems, and we will collaborate on using PCT to develop diagnostic procedures for, and models of, various neurological impairments. And as I mentioned during some of the earlier discussions on the net, we will try to entice a physiologist or two into looking for evidence of control system organization in the nervous systems of "simple" creatures.

On that last topic, I do not think most people are aware of how decidedly skewed the physiological literature is, due to the nearly universal assumption that "simple" creatures act as tropistic, or reflexive, or instinctual, systems. In every such case, the presumption is that the creatures act open loop, then experimental procedures are devised in which evidence for open loop actions appears, then the search for physiological mechanisms unsues and often concludes with what appears to be support for open loop processes. These presumptions and procedures are nearly universal. Consequently, much of the evidence in the literature deserves re-examination, if for no other reason than to effectively rule out the possibility that genuine feedback control processes are at work at the level of the overt behavior of the organism.

The existence of a massive research literature in no way assures that there exists a body of data that are relevant to the task of ruling in favor of, or against, either linear or circular causality. The extensive literature on the "conditioning" of animals, and people, provides a case in point. Out of a large library of work on operant conditioning of animals, there are probably fewer than ten articles that even report the right kinds of data for one to use in comparative modeling. And only a few studies have allowed animals to control their environments, and the authors openly asserted that the animals indeed exercised such control.

How, then, can there be any assurance that research into the physiological correlates of actions of animals is even near to being on track? For example, are there any data on the activity of sensory systems in moths, during their free-falls after first hearing the echo-ranging sounds of bats? Or of the activity of those systems after the moth is on the ground? Or are there close accounts of the activity of falling and grounded moths? Are their actions really identical in all cases, or only "sort of close," as is the case most often when people swear that two events, isolated in time and space, are "identical," but actually are anything but?

Tom Bourbon   <TBourbon@SFAustin.BitNet>
Dept. of Psychology
Stephen F. Austin State Univ.
Nacogdoches, TX 75962   Ph. (409)568-4402

Date: Wed Mar 25, 1992 4:30 am PST
Subject: "Purpose" vs. Adaptation to Purpose

From Greg Williams (920324)

Subject: Re: Malcolm's automata

>Chris Malcolm (920324)

>There is a definition of purpose and disturbance available
>in which the robots can clearly be seen to be achieving their purpose
>despite disturbances.

I guess you would think it strange to attribute "purpose" to a bullet, rather than to the person firing a gun (sometimes -- non-purposeful firing of guns is possible!). The bullet was designed so as to not be affected much by the typical disturbances it might encounter along its (ballistic!) trajectory. I was a bit sloppy in my post yesterday when I hypothesized that you ELIMINATED environmental disturbances affecting your automata -- that would be impossible, of course, in the real (even laboratory) world (though not in computer simulations, as Bill pointed out). As you say, you only needed to reduce the effects of certain disturbances sufficiently for... what? NOT for your automata to "achieve their purposes," but rather TO SHOW RELIABLE ADAPTATION TO YOUR OWN PURPOSES. I was careful in my post yesterday to refer to the accomplishments of your automata as "adapted" actions, since I wouldn't want to ascribe "purpose" to those automata, as I wouldn't want to ascribe it to a bullet.

>They achieve their purpose always, without perceptual control, provided
>the environment is kept within the design limits.

I would say: "They achieve their ADAPTATION TO YOUR PURPOSES always, ..." Why is it so important to distinguish between purpose and adaptation to (I nearly said "other" here, but your automata don't count as agents) agents' purposes? To avoid the muddles which you raise below! To avoid such muddles, one must reserve "purpose" for perceptual controllers with internal reference levels which don't (just, although they might also) pre-compensate for disturbances, but actively OPPOSE disturbances.

>Of course their purposes can be frustrated by disturbances they can't cope
>with, such as setting them on fire or pouring a bucket of sand over them, but
>the same kind of restrictions apply to perceptual control.

Both types of systems do have limits to the disturbances they can handle. But up to those limits, they counter disturbances in two quite different ways: passive vs. active. The Test, considered as a Test for Purpose, can be tightened up by asking whether the actions involved in maintaining some variable nearly constant involve actively "mirroring" the disturbances. If there is no "mirroring," just "going with" the (within-design-tolerance) disturbances, then there is no purpose -- you are looking at a fancy bullet. But the Test is not infallible, and you ultimately need to look inside the system, at least in some cases, to see a PCT organization with a reference signal, and thus become fully convinced that it is a purposive system.

>Given a robot, an environment variable within limits, and a purpose achieved
>despite these variations, why should the details of the process by which I
>arrived at the design matter?

If your givens are to be accepted, you had better be explicit that the "purpose achieved" is the purpose of the (PC-organized) designer. If the design was non-purposive, then one of your givens disappears. Suppose you purposefully design a balancing automaton which is successfully adapted to your purposes; you can legitimately say that there is a purpose achieved -- YOUR purpose. Now suppose you throw a stick and it "happens to" balance over another stick; you may not legitimately say that any purpose has been achieved.

>Are you arguing that my automata can only borrow my own purposes? That sounds
>horribly like the traditional "magic" (intentional, subjective, etc.) view of
>purpose which I thought perceptual control was supposed to be able to rescue
>us from?

No, as you can see from the above. The "magic" comes in when you fail to distinguish purpose from adaptation and PCT-type purpose from other sorts of "purpose."

>So in talking about the purposes of simple mindless creatures (supposing for
>the sake of argument that dung beetles are mindless) we find we are implicitly
>talking about the "perceptions" of evolution. Really? What happened to the
>idea of _objectively_ discovering purpose by _experimentally_ discovering the
>controlled perceptions?

You can get yourself into all sorts of difficulties by talking about "purposes" of feedforward devices. I've already addressed above the experimental approach to finding behavioral purpose in any organism, "mindless" or not. However (see my post yesterday on nonartifactual automata), To go WAY out on a limb, I think that the evolutionary process might indeed have an organization somewhat analogous to organismic PC-purpose. Evolution happens because genetically variation gives rise to variation in reproductive success. One could PERHAPS speak profitably of a (metaphorical) "perceptual" feedback in evolution connecting the actions of an organism (modified by "disturbances") with the "perceived" outcome, reproductive success. Break the connection between the actions and reproductive success, and there is no evolution. But I don't understand what an evolutionary "reference signal" might be. Regardless, one can simply recognize the evolutionary process as a surrogate agent to which either purposive or non-purposive activities of organisms can be adapted.

>... you can discover the purpose of my "senseless" robots by seeing what
>happens when you change things. You do not need to ask the designer, God, or
>evolution.

You can discover whether or not their actions are adapted to an agent's purpose only by considering how they came to be. For nonartifactual feedforward devices, you can discover whether or not their actions are adapted evolutionarily (as ballooning of spiders presumably is) or are not (as dropping of stones off cliffs presumably isn't).

Greg

Date: Wed Mar 25, 1992 10:58 am PST
Subject: Re: Malcolm's automata

[From Rick Marken (920325)]

I just noticed this comment latest from Chris Malcolm's last reply to Greg Williams:

>As it happens I didn't trim the environment to fit the automaton, I
>improved the automaton to handle the environment. But this is simply an
>implementation detail. Given a robot, an environment variable within
>limits, and a purpose achieved despite these variations, why should the
>details of the process by which I arrived at the design matter?

I take it that this means that we are not going to get the details on the design of Chris' non-sensory, purposeful robots. A later statement of Chris's suggests that Chris considers the behavior of these robots to be purposeful -- just a different kind of purpose than that exhibited by perceptual control systems. For example, Chris says:

>I applaud that view. My purpose is to argue for its extension, i.e.,
>there are more purposes than those achieved by perceptual control, but
>they can still be discovered by observation and experiment of a similar
>kind, i.e., you can discover the purpose of my "senseless" robots by
>seeing what happens when you change things.

I think, then, Chris agrees that the kind of purpose exhibited by his robots is NOT control. It is more like the purpose exhibited by a pendulum that returns to its resting state after a transient disturbance. If you want to call this purpose (and many people besides Chris have done this) then what can I say? That's your choice; but it seems to eliminate much of the value that language might have for communication. Sort of like Humpty-Dumpty science. I must hand it to Chris, though. At least he admits there are other kinds of purpose besides the kind exhibited by his robots; he admits that there is the PCT kind too. This is more than I can say for others who also seem to believe in Chris's robot's type of purpose; these are the "coordinative structure, dynamic equilibrium, point attractor" types who believe that these open loop systems exhibit the ONLY kind of purpose. I don't think control theory can help people who want to see ordinary physical processes (like raindrops converging into a common stream or pistons moving up and down in a cylandar) as evidence of purpose (on the part of the raindrops or the pistons). Control theory is about control; if the phenomenon is not control (as we quantitatively define it) then, call it what you want, PCT is not applicable to or threatened by its existence.

Regards   Rick

Date: Wed Mar 25, 1992 2:10 pm PST
Subject: Open loops, closed loops and HCI

[From Rick Marken (920325 13:00)]

Martin Taylor (920323 16:30) says:

>But, in the field of HCI, we see numerous situations in which the interface
>designer has precluded the closing of the loop, at least forced it to a higher
>level.

> I think one of the main
>benefits of direct manipulation interfaces is that they provide continuous
>possibilities for the loop to be closed. The user can see all the time
>how the actions are affecting the perceived state of the machine. This
>aspect of the benefits of direct manipulation is never (I think) discussed.
>Instead, the idea is put about that such interfaces eliminate the possibility
>of error, which they don't.

Maybe we could start a thread here that is actually relevant to my work. I think there are many issues in the field of Human-Computer Interface (HCI) that are relevant to PCT. Martin mentioned two in his comments above : 1) a person working on a computer is in a (hopefully negative) feedback control loop; the (intermediate level) controlled variables are on the screen. 2) "error reduction" is one of the big concerns in the field (others include "efficiency", "usability", "safety"); obviously, HCI people typically use the term "error" to describe descrepencies between what they think a result should be and the result being produced by the operator. Thus, the "errors" are experienced by the HCI engineer, not necessarily by the operator (like all the spelling "errors" that I happily and confidently produce in my posts; usually they are only errors for you, not, unfortunately, for me.

There are other interesting topics related to HCI that might also be discussed. One that I think is interesting is the fact that theories of HCI almost always recognize the closed loop nature of the HCI process (the best example is Don Norman's "User Centered" (I think that's what it's called) model of HCI) but fail to understand that this means that the operator is involved in control of perception and that output (keypresses, commanding, etc) is a function of disturbances and the feedback function, not input (basically what is on the screen).

I don't know where to begin with a discussion of HCI -- or what we might get out of it -- but I'd love to start such a thread.

I will say that I agree with Martin's basic point about "direct manipulation" type interfaces; the connection between output and input may be more explicit than in other interfaces. For example, clicking and dragging can be used to move a file icon to the trash. Once you learn to drag icons and how you can do things with them by dragging you can get a lot of results by just clicking and dragging. In the command oriented approach you have to remember the letter sequences that produce the same result; so more of the "how" part of getting things done is stored in memory (not easy for some of us -- like me).

My main HCI project here at work right know is assisting in the development of standards for satellite control HCI. The goal is to develop some conventional ways of doing satellite control tasks that could be adopted by all the agencies involved in satellite control -- Military, NASA, possibly International Agencies. The standards should increase "interoperability" -- making it possible for people trained to operate satellite system A to move, with minimum retraining, to satellite system B. They should also reduce procurement costs because contractors would not have to design a completely new interface every time there is a block upgrade to some satellite system.

Anyway, that is one of my main interests at the moment -- HCI Standards. Somehow I think PCT can help to at least organize the problem. For example, PCT suggests two places where standards might help 1) in the feedback function relating output to input (standard ways of affecting the display) and 2) in the displays themselves so that potential controlled variables are represented in the same way.

Any thoughts on HCI standardization would be welcome (including thoughts like "what a waste of time"). I would be particularly interested in ideas about criteria for selecting standards, possibly ways of testing and evaluating proposed standards, etc. When are standards arbitrary (like "steering wheel on the left") and when not (if steering wheel on left then traffic on right)?

Gary Cziko -- what is the CSGNet policy regarding actually doing work over the net? I guess it should be OK -- but it seems like it could actually make work fun. And you know me; I always think something must be wrong if I'm having a good time.

Best regards Rick

Date: Wed Mar 25, 1992 2:22 pm PST
Subject: uudecode problems

[Wayne Hershberger 920325]

Bill Powers and/or Gary Cziko

I am having trouble getting a working copy of the new versions of Demo1 and Demo2. I got the ASCII files over bitnet. I combined them and uudecoded them getting a single file in each case, dem1a.exe, and dem2a.exe. When I attempt to execute either of these files, nothing happens, no error messages, nor screen cursor, nothing; it just bombs, I have to reboot the computer. Have other people managed to get the uudecoded files to work? Have you head? Do you have any advice for me?

Frustrated, Wayne   tj0wah1@niu.bitnet

Date: Wed Mar 25, 1992 3:00 pm PST
Subject: Cockroach escape reorganization

[from Gary Cziko 920325.1350]

In reply to my (920324.1320) saying:

>>How can the cockroach change his behavior adaptively in this situation if
>>it doesn't care where its going (e.g., make sure ground speed is faster
>>than air speed, that means I going downwind, away from the air disturber)?
>>Please, somebody explain to me how this can happen if the cockroach's
>>escape behavior is purely ballistic.

Greg Williams (920324-3) says:

>The precalibrated settings for ballistic escape might be reset on the basis of
>wind direction/intensity AFTER the escape was completed. No feedback control
>would be needed during the escape for this to happen.

Yes, I suppose that this is possible, but I don't like it very much. If this is what's happening, the cockroach has to sense something after the escape routine is finished and then realize "Hey, this isn't where I was supposed to end up. I better try a new twist to my escape routine the next time I feel that puff of air. I hope I remember. And I hope that what I try doesn't put me in a worse position than I'm in now (e.g., under the cockroach stomper's boot heel)."

It seems to me that the reorganization would seem a lot easier and faster if the cockroach is getting sensory feedback WHILE it is escaping. Then he could reorganize successfully (at least to some degree) in one trial.

If we saw the cockroach with the newly waxed cerci getting closer to the right direction of escape toward the END of the escape move, this would be good evidence that the move is under feedback control of some type. Here is where some good cockroach data would be of use. I think I'll send this to Randy Beer and see what he has to say.--Gary

P.S. That reminds to remind CSGnetters once again that Randy Beer is NOT on CSGnet, although he sends his replies to Powers, Marken, et al. to CSGnet. If you want to have your message sent to Beer you have to add his address after the CSGnet address. This is beer@cthulu.ces.cwru.edu (Tom

Bourbon take special note).

Gary A. Cziko

Date: Wed Mar 25, 1992 3:05 pm PST
Subject: Re: "edge of chaos"

[Martin Taylor 920325 17:20]
(Jeff Dooley 920324.1900)

It's interesting that you should bring up the "edge of chaos" here. I tried to do that when I first joined this group, and got shot down for it. I am hoping that I can bring the mutual understanding of the group to a position where it can accept it as a natural part of PCT, which I think it is. But I have been holding fire for a year or so, now, because I think many of the group will misinterpret it.

I started to write a paper based on entirely different foundations about three years ago, entitled "Thoughts on the edge of chaos", which argued that ANY thinking machine that would be said to be "intelligent" would have to be operating on the edge of chaos, and moreover would have to use catastrophe functions that were the "elements" of the "critical landslide" (to use the metaphor chosen by whoever wrote the Scientific American article last year). The catastrophe functions correspond to "categories" in perception, and are the only things that permit psuedo-logical operation in physical systems. (In computers, the catastrophe functions are built-in at the lowest level, and do not have a fractal structure, which makes real "artificial intelligence" very hard to achieve).

Anyway, I never finished the draft paper, because it was intended for Behavioral and Brain Sciences, and Freeman got his chaos paper in first.

For PCT people, the key point will be that there must be same-level connections in the hierarchy, but we are yet a long way from having developed an agreed substrate for developing that understanding. We'll get there, I hope, before another year passes.

Martin

Date: Wed Mar 25, 1992 4:28 pm PST
Subject: Incommensurability; Tight Links

[from Gary Cziko 920325.1720]

Jeff Dooley (920323.1100) says:

>Finally, a footnote on different ways of "seeing" the world:
>The two paradigms of S-R, open-loop, etc, and PCT, closed
>loop, etc. seem incommensurable in a Kuhnian sense. Workers
>in the different traditions really do seem to think, speak,
>operate in, or at least "see," different worlds. This could
>be serious, since theoretical development, experimental
>design, and other activities influencing the direction of
>science are informed, one could say, as a function of which
>world one belonged to. Could incommensurability in this

>sense help explain the apparent fact that S-R types and
>PCTers seem to just talk past one another?

You may call it incommensurability, but perhaps it can be understood as controlling for two different types of perception--one which sees one-way, cause-effect relationships accounting for behavior which may LOOK purposive (but isn't), and the other which sees closed-loop relatioships accounting for behavior which IS purposeful. Each camp resists disturbances to its view, until too much error remains and reorganization takes place.

In these types of paradigm shifts, I find it intriguing to consider the direction of movement. I once had a very S-R perspective on behavior but I can now see all kinds of problems with it and can't see how I would ever return to it. I'm sure there must be others on CSGnet like this. But are there any individuals who ever really understood PCT and then defected for S-R or other incompatible viewpoints? If not, this would lead me to suspect that the PCT is indeed an improvement in that it accounts for apparently S-R phenomena as well as phenomena that S-R cannot explain.

But in the history of science, does the direction of movement of individuals correlate with the eventual widespread acceptance of new theories and paradigms? If there are people moving form theory A to theory B but not the converse does this mean that theory B eventually replaces theory B? Perhaps Dennis Delprato can fill us in here. We haven't heard from him in a while. In the meantime perhaps Bill Powers can let us know if he knows of anyone who had a good understanding of PCT who nonetheless abandoned it.

===

Bill Powers (920324.1100) in his excellent response to Randy Beer said:

>So this kind of system [open loop]
>can work only in an unreal or protected environment in which such
>disturbances can't happen -- or else if the variable is so tightly
>coupled to the plant's output that it can't vary from the state fixed by
>that output even when disturbances are present.

Could you explain to me what "so tightly coupled" means? You seem to contrast it with a situation where disturbances can't happen, so it seems you mean something more than just that the disturbances are too weak to have an effect.--Gary

Date: Wed Mar 25, 1992 5:01 pm PST
Subject: Edge of chaos; purposive behavior

[From Bill Powers (920325.1100)]

Jeff Dooley (920324) --

Astute of you to pick up Kauffman's idea. I recognize what he is talking about from experiments with reorganization, the E. coli models. Here are some impressions I picked up, generalizations without benefit of systematic verification.

There appears to be an optimum loop gain in a reorganizing system. If the gain's too low, even large errors don't produce very frequent "tumbles," so progress toward lower error is very slow. But if it's too high, favorable directions of movement still produce a relatively short interval between changes, so there isn't a chance for much progress toward the goal even when the correct direction happens to result. Somewhere in between there's an optimum relationship between the error signal and the tumbling rate. Intuitively, I can see that reorganization has to allow enough time for evaluating the consequences of the new "direction of movement" and to take advantage of favorable directions.

I don't think the edge of chaos is as "razor thin" as you suggest. There's probably a rather broad range over which the goal would be reached in time for survival. Of course Kauffman seems to be thinking in terms of the scope of the reorganizing changes rather than their frequency, but even so I would think that there would be some leeway -- E. Coli's method of travel seems to work just as well with three-dimensional tumbles as with 1-dimensional ones, so the success isn't extremely sensitive to the number of degrees of freedom simultaneously being changed.

I'll let Rick's good answer to your previous post stand in for mine.

Tom Bourbon (920325) --

You will be the envy of all pure CSGers. Some advice from my days long ago at the VA Research Hospital, where I was supposed to have half of my time to devote to control theory: learn quickly how to say no. When people see that you actually know something, they will try to drag you into their worst-conceived projects to rescue them. Let them sink. If you say yes, as I did far too often, you'll end up wasting all your time on trivial projects that need total revision to work. To reveal my deepest prejudices based on too small a data base, my strongest advice is to avoid surgeons. All of them I met had incredibly inflated concepts of the brilliance of their own ideas and they assumed that everyone would just go along. I think it has something to do with medical school.
On the other hand, maybe I just got a bad batch. But try to stick to control theory if they'll let you.

You comments about baby spiders are excellent. To build on that, I think we also have to remember that the chances of arriving in one evolutionary jump at the production of baby spiders who disperse this way is zero. Before they could disperse this way, they had to be able to spin silk at an early age, climb plants and trees to high places, learn to jump off instead of hang on, and of course do all the things that are needed to locomote up irregular surfaces in a systematic way, counter to gravity. The few examples of open-loop behavior we see -- and apparently the only kinds of behavior that traditional scientists are capable of recognizing when they see them -- are supported by a vast hierarchical structure of control systems. If there really are any open-loop behaviors, their components are all control behaviors, and as you say, they will remain only until long-term environmental disturbances make them counterproductive again.
Evolution might throw up an occasional open-loop reponse to a stimulus, but its chances of survival in the company of other variants that do the same thing by control -- in a disturbance-resistant way -- are pretty poor. It's sad, but when the only behavior you can recognize is a response to a stimulus, you're going to miss practically all of what is going on.

As to control vs. noncontrol, we're up against history. The only widely known alternative to the "scientific" concept of purpose (outcome) is the metaphysical concept. To say that organisms have INNER purposes is, in the view of most scientists, to classify yourself as a metaphysician. Somehow scientists have to learn that there's an alternative both to the current scientific view and to metaphysics. Not many of them go far enough into control theory to realize that they're looking at something radically new. What we need is come sort of quantum tunneling argument.

Greg Williams (920324): a very astute comment.

I'm pleased that so many different people are independently coming up with the correct arguments against Beer and (sorry Chris) Malcolm. The CT concept of purpose does NOT mean just a consequence of an action, but an INTENDED consequence. In all systems without feedback, it is only the human designer and user who can see whether the consequence of a "purposeful" (useful) action is the desired one, and take action to correct the result if it's not. You can say that the purpose of a lawnmower is to cut grass, but if you were to go right now to the place where you keep your lawnmower, I'll bet any amount it would just be sitting there, not cutting grass or accomplishing any other useful purpose. The human designer and operator of these so-called "purposive" open-loop systems seems to be completely invisible. We have to keep calling attention to that man behind the curtain madly working the levers while the Great Oz roars.

Best to all

Bill P.

Date: Wed Mar 25, 1992 6:21 pm PST
Subject: Re: Open loops, closed loops and HCI

[Martin Taylor 920325 21:00]
(Rick Marken 920325 13:00)
>
>
>Gary Cziko -- what is the CSGNet policy regarding actually doing work
>over the net? I guess it should be OK -- but it seems like it could
>actually make work fun. And you know me; I always think something must
>be wrong if I'm having a good time.

Gary, I hope that you would answer "Work on the net is OK, if it illuminates or makes concrete aspects of PCT, but not if it is just for the benefit of the dyad or small group concerned." As an interested party, I'd like to pursue Rick's idea of looking at real concrete HCI problems, but I can do that with him off-line if the questions are of insufficiently general interest.

Rick, I like the idea, and I think that discussions of HCI might illuminate the ever more technical discussions of language, considering the computer as having characteristics that, from a PCT viewpoint, are intermediate between a human and a screwdriver.

It's a funny reference signal that generates error unless you detect error. I LIKE having a good time as part of my work!

Martin Taylor

Date: Wed Mar 25, 1992 7:07 pm PST
Subject: No Evolutionary Reference Signal

From Greg Williams (920325-2)

Before the slender limb supporting my half-baked ideas on evolution cracks even more (how's that for a mixed(-up) metaphor?), I want to recant (at least for now -- still baking) the possibility of some sort of evolutionary reference signal. It looks like good old "blind variation and selective retention" suffices as a surrogate for a PC-organized agent in generating adaptations.

I should have realized earlier that if intraorganismic reorganization can generate novel adaptations, either PC-organized or "feedforward" -- which I do think is the case -- then an analogous interorganismic reorganization can also be expected to generate novel adaptations.

It's as if a designer were doing what Chris Malcolm said he did to make his non-sensing robots do what he wanted, except by running many trials in parallel on differently organized robots, throwing away the ones that don't "work" and increasing the number of ones which do "work," where "work" means "would be highly successful at reproducing, were that possible" (or something like that -- I don't know enough to be really precise about it -- maybe that unpacks to "would be able to produce a lot of offspring which reproduce"?), meanwhile tinkering a little IN A NON-GOAL-DIRECTED WAY with the innards (mutations and crossings-over). Of course, in real evolution, since some of the individual organisms actually CAN reproduce, the offspring don't actually get thrown away or built by a designer, they just do or don't get born (or do or don't survive to reproduce?). No agent-designer is needed for the selective retention part of evolution IF reproduction generally means replication of parental behavioral organization, AND some behavioral organizations reliably correlate (over several generations) with high reproductive success.

Greg

Date: Thu Mar 26, 1992 6:37 am PST
Subject: Open loop vs. closed loop

[From Bill Powers (920325.2000)]

There's an important point about the open-loop, closed-loop argument that we've been missing; I really feel slow in not catching it until now, while reading Rick's comments about SR systems changing into closed-loop systems and vice versa depending on the circumstances.

THERE ARE NO CLOSED-LOOP BEHAVING SYSTEMS (except those with negative feedback internal to their nervous systems). The contrast we should be making is between closed-loop and open-loop _situations_. If, in the external world, there are connections such that an action by an organism has an immediate effect on the relevant input to the organism ("immediate" being defined in terms of the speed of action of the system), then the situation is closed-loop. If the action has no immediate effect on the relevant input, the situation is open-loop. You don't have to understand the organization of the behaving system to distinguish open-loop from closed-loop situations.

To say that an organism is organized as a control system with respect to some specific stimulus is to say three things: first, that it subtracts its input from an internally generated reference signal (or vice versa depending on the external situation), to establish the effective zero point of the input. Second, that its action is based on departures of the input (as analogized in a perceptual signal) from the reference level (normally specified as a reference signal). Third, that in the environment in which the behaving system evolved, there is a strong effect of the output on the input, IF THE NORMAL LINK EXISTS. The sense of this effect will be to create feedback, the sign of the internal comparison or output process being chosen for negative, rather than positive, feedback.

If a control system became organized in a specific environment so that a strong external feedback link normally exists, it will behave in all respects as a control system, controlling its own input information. The environment, however, can change.

For example, suppose you're driving down a twisting mountain road and your car suffers a complete electrical failure, ignition and all. Where you had been negotiating the curves effortlessly, you now suddenly start making huge steering efforts, almost more than you can produce, because _your power steering is gone_. In the environment that normally connects your steering efforts to the perceived position of the car on the road, the feedback link has suddenly become much weaker; a given torque applied to the steering wheel now has far less effect on the lateral motions of the car.

Your internal system still has the same properties it had before. A given perceptual error in position of the car still leads to the same increase in steering effort, in the appropriate direction, as before. But the part of the loop gain contributed by the external part of the loop has suddenly dropped by a large factor, so the total loop gain has decreased to a small fraction of its former value. As a result the error increases greatly, causing greatly increased steering efforts. But your steering efforts, large as they now are, control the perceived path of the car far less effectively.

And what if the steering mechanism broke completely? You would still be organized the same way inside, appropriately for controlling your perceptions in a normal environment. But now the car would deviate greatly from the path you want to see; the error signal would become enormous; your output efforts would swing back and forth wildly between their maximum limits in the final moments before the careening car finally went off the road.

In this last scenario, you are still organized as a control system but you're no longer in a closed-loop situation. Because the external feedback link is gone, you are operating open-loop. The extreme actions show that the error signal is far out of its normal range, and in fact indicate that you have lost control.

Any organism that has evolved to control its own perceptions in a particular environment can find itself in an open-loop situation. This can happen not just through losing the feedback connection, but through encountering such a large disturbance that your efforts to oppose it saturate. Once you're producing maximum output to oppose a disturbance, any further increase in that disturbance affects your input without opposition. The loop is broken because now changes in the disturbance, which cause changes in the perception and the error signal, no longer produce matched opposing changes of output. So the situation has become open loop even though you're exerting the maximum possible effort to maintain control.

In stimulus-response experiments, most often the applied stimulus is really just a disturbance of some other input variable undetected by the experimenter, a disturbance which the test organism can successfully cancel by altering its actions. This leads, as I have mentioned before, to illusory stimulus-response laws that really reveal only environmental properties.

But in some experiments, the experimenter gets hold of the actual input that's being controlled. Rather than applying disturbing _influences_ to that input, which the animal can counteract, the experimenter puts his own vastly more powerful control loop to work and forces the input to change regardless of the animal's efforts. This is called "varying the independent variable." Doinhg so effectively puts the animal in an open-loop situation, because its actions no longer affect its inputs. Now what you see are strong reactions to the input, because the changes of input directly affect the error signal, which is usually highly amplified to produce the output. If the loop were closed, this high amplification would not create large outputs; it would just keep the error small, the outputs being only what is necessary to counteract normal disturbances. With the loop open, however, this amplification creates extremes of output. The system is being operated in a highly abnormal condition.

Footnote: I will never forget reading of an experiment in which the researchers were trying to control for every possible interference and get a reliable response to a stimulus out of a rat. They strapped the rat into a narrow box and sewed its eyelids open so it couldn't avoid seeing the stimulus light. They got the same response (from a leg, I think), after conditioning, something like 80 per cent of the time, and gave up. They were studying this rat in a totally open-loop situation -- how they thought they would get any data about normal behavior escaped me then and escapes me now.

When an organism that normally acts as a control system experiences an open-loop situation, it becomes hyper-responsive to changes in its input because those changes are no longer counteracted; they show up directly as error signals. There will be hyperresponsivity to changes in reference signals, too, because the reference signal also shows up directly as a change in error signal. Normally the perceptual signal would immediately catch up and the error would be kept from getting large. But with the loop open, the perceptual signal no longer changes because of the output, and the error signal remains large.

If the loop is opened by denervation, usually only lower levels of control are affected. Given time, the higher level systems that normally use those lower-level systems as means of action will reorganize to compensate for the overresponse of the lower systems now running open-loop. The initial instability caused by too high a loop gain gradually disappears as reorganization lowers the gain in the superordinate control systems. Nothing can be done about lower-level feedback dynamics that serve to stabilize limbs; the higher systems are too slow to compensate fully for dynamical effects. But higher-level control of a sort will be restored. This is what is meant when researchers who use denervation methods say that denervation shows that feedback is not necessary for "normal" behavior.

While there are no closed-loop systems, there are open-loop systems, systems containing no provision for their actions to affect sensitive sensory inputs, so that all situations are open loop.

Organisms have evolved to take advantage of the fact that their outputs affect their own sensory inputs. In fact they have evolved elaborate sensory systems specifically designed to detect the effects of essentially every possible action, external and internal to the body.

But what of systems so organized that there are no such effects in any environment? How do these systems have to be organized in order to have reliable objective effects on their environments? It is possible in principle for such systems to evolve, even among organisms, strictly on the basis that the objective effects of their actions affect their survival to the age of reproduction. The question is, what properties must evolve so that the resulting actions will counteract external influences that interfere with surviving to reproduce?

First, the actions must be protected from external disturbances that could change their effects, or else must be produced by such a massive mechanism that normal disturbances are incapable of altering the effects, or else must somehow be compensated without feedback (see below).

Second, the actions must be produced in a uniform way, so that the output calibration of the system in terms of external outcomes will never change enough to alter the critical effector outputs or their objective consequences.

Third, if the actions are based on sensory inputs, the calibration of the sensory inputs must also remain absolutely stable, so that the same external situation will always result in the same effective stimulus.

Fourth, if significant disturbances remain possible, then the sensory system must detect each separate possible cause of a disturbance, convert its state reliably into a calculated effect on the outcome, and inject a compensating stimulus into the system that opposes the effect on the outcome, adjusted to include the behaving system's own properties, all dynamical effects, all nonlinearities, all changes in the relationship of the potential disturbance to the outcome, and all changes in the link between effectors and outcome.

If this set of requirements doesn't seem beyond meeting, there are more. In most behaviors, the critical outcome doesn't depend directly on the effector output, but on other variables that depend, often loosely, on the effector output. When we consider locomotion, the arrival of an organism at a particular place, or even the placement of its limbs in a particular orientation, results from the application of muscle forces to limbs, and the subsequence effects of limb forces on other objects. To go from forces to positions requires two time integrations, nonlinear ones when jointed limbs are involved. Time integrations are notoriously sensitive -- hypersensitive -- not only to initial conditions, but to very small errors of computation. In living systems we have to add the fact that even an accurately computed output can't be translated accurately into output forces by real neurons and muscles; these inaccuracies, too, contribute to the error in the final integrated result. While a simple brief movement might come somewhere near the necessary result, a long series of movements occurring serially, such as walking across a room to a source of food, would begin each new movement with all the accumulated errors of the previous movements. These errors would remain undetected, because of the absence of feedback information. The only way of detecting failure would be to fail and die.

I think the only reasonable conclusion is that no behavior of even moderate complexity and short duration can be counted on to produce a reliable effect in an open-loop situation. Given the wonders of electronics and incredibly accurate mechanical constructions, and environments free of unpredictable disturbances, open-loop systems can produce reliable results far into the future, requiring only occasional mid-course corrections. But this is not even remotely possible for living organisms whose input sensitivites vary, whose muscles fatigue, and which live in environments where disturbances are ubiquitous and mostly hidden from the senses.

So open-loop systems can exist. Even control systems, in an open-loop situation, will behave like open-loop systems: inaccurately and unreliably, if they are living systems rather than marvels of mechanical engineering and stable precision electronics. In the competition for survival, open-loop systems which can't detect the consequences of their actions directly, while they are being brought about, don't stand a chance when their competitors are control systems equipped to sense the outcomes of their own actions and control them.

Gary Cziko asked what I meant by "tight coupling" in a recent reply to Randy Beer. In one case presented by Beer, there was an output that was sensed and in fact controlled, but the real "controlled" variable was some other effect of the output, not sensed. I said that if this other variable were tightly coupled to the controlled output, it might well be unaffected by disturbances, and so effectively be controlled by the system.

Suppose the controlled output of the "plant" is a motor, and that the controlled aspect of the motor output is rotational velocity. A tachometer will provide a signal representing angular velocity. This signal will be compared with the input signal specifying desired angular velocity, and the error will be amplified and used to drive the motor. The result will be that disturbances like varying loads on the motor will have little effect on the speed of the motor. With a sensitive control system, the effects of varying loads can be made undetectable.

Now suppose that the real controlled variable is to be the rotational velocity of a wheel. If the motor has a thick short shaft on which the wheel is directly mounted, then applying braking forces to the wheel will not be able to slow the wheel, because slowing the wheel would entail slowing the motor, and the tachometer signal would drop very slightly, raising the drive to the motor enough to prevent any significant drop in speed. This is "tight coupling."

On the other hand, suppose there is a fluid-drive transmission between the motor and the wheel whose speed is to be "controlled" in this way. If a braking or accelerating torque is applied to the wheel, the wheel will begin to turn slower or faster than the motor. The motor itself will continue to turn at a constant speed, because its speed is controlled. But the wheel will slip or advance relative to the motor, because of the "loose coupling" through which it is driven. The fluid drive mechanism will apply some corrective force, but the speed of the wheel will not be controlled nearly as well as that of the motor.

If we now imagine the shaft connecting the motor to the wheel as a quarter-inch diameter steel rod 100 feet long, it's clear that the wheel's angular position can easily be disturbed relative to the motor shaft's angular position. We would probably stop the wheel by hand, briefly, while the motor put a twist into the long shaft. The motor itself would continue to spin at an inexorably-controlled speed, but the wheel could easily be disturbed in angular position and momentary speed over a wide range of variation. That would be very loose coupling.

In both of the loose-coupling conditions, accurate control could be restored (mostly) by moving the tachometer to the position of the wheel, so it measures directly the speed of the wheel. In the case of the long shaft, some dynamical filtering would be needed for stable control, but in the end the wheel would resist braking or accelerating torques with great precision (even if it might respond a little more slowly to sudden torques than the tightly-coupled wheel would).

Of course now the motor shaft speed would become uncontrolled. A braking force applied to the wheel would cause the motor to speed up -- permanently in the case of the fluid-drive coupling, temporarily for the long shaft, while the motor quickly wound up the shaft to create the necessary opposing torque.

In "control-of-output" interpretations, it's always assumed that the controlled output variable is tightly, nay rigidly, coupled to the position where the feedback sensor is located.

Note that there is no way for an open-loop system to control a motor's shaft speed without sensing the braking or accelerating forces applied to the shaft by outside agencies. And then the control (actually, compensation, which is not control) will only be as good as the calibration, linearity, and constancy of the sensors and the motor's response to driving signals.

If any CSGer wants to have the kind of nauseating experience that Rick likes to celebrate, read the opening main article in Science for 20 March, 1992.

From Mary Powers:
[from Mary Powers]

Gary Cziko (920325) asks if Bill knows of anyone who understood PCT well who has abandoned it.

Well, these are the people he ISN'T in touch with - with his full plate it's hard to notice what's missing (or who).

I would suppose that a reasonable group to look into would be people who came to at least two meetings of the CSG but are not current CSG members. If you are really interested, I could probably dig out the lists of attendees at CSG meetings (good motive to do some filing I've been avoiding). What sort of questions would you ask the drop-outs?

Mary P.

Best to all from both of us

Bill P.

Date: Thu Mar 26, 1992 6:47 am PST
Subject: open loops

[Avery Andrews (920326.1750443)]

(Bill Powers (920325.2000))

>If the action has no immediate effect on the
>relevant input, the situation is open-loop. You don't have to understand
>the organization of the behaving system to distinguish open-loop from
>closed-loop situations.

But shouldn't we distinguish cases where the action (output = O)'s effect on the input (I) flows thru to effect further changes in O from those where it doesn't? E.g., in the infamous plummeting moth, it is perhaps the case that the change in position of the moth caused by folding up (O) causes some change to I, but there will not typically be any resulting further change to O (assuming what Randy Beer says about moths being able to hear bats a long way off), and even if there were (a slightly tighter or laxer hunker), they wouldn't have much in the way of a further effect upon I).

It seems to me that the case where O has no effect on I, and the case where it does, but this has no further effect on O, are similar, and are the ones that can do useful work only under extremely limited circumstances (which appear to be satisfied in the case of the moths, since they wouldn't be flying around in the first place in weather conditions where falling didn't cause plummeting.

I have no problem at all with the belief that (sub-)systems that are supposed to run open loop (my sense) are pretty rare, but I think it's quite important to accept them without a fuss if they stand up to careful scrutiny. If you don't, people are likely to get the idea that PCT is some kind of religion rather than an actual insight into what is usually going on with living things.

Avery.Andrews@anu.edu.au

Date: Thu Mar 26, 1992 6:50 am PST
Subject: Re: uudecode problems

>I am having trouble getting a working copy of the new versions of
>Demo1 and Demo2. I got the ASCII files over bitnet. I combined
>them and uudecoded them getting a single file in each case,
>dem1a.exe, and dem2a.exe. When I attempt to execute either of
>these files, nothing happens, no error messages, nor screen cursor,
>nothing; it just bombs, I have to reboot the computer. Have other

>people managed to get the uudecoded files to work? Have you head?
>Do you have any advice for me?

Please let me know exactly how you are decoding them. There are many enhanced decoders available, but if you are using the standard Unix uudecode you have to concatenate the parts and remove the headers. The decoder uue.c on biome doesn't require this editing.

Bill
--
Bill Silvert at the Bedford Institute of Oceanography
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2
InterNet Address: bill@biome.bio.ns.ca

Date: Thu Mar 26, 1992 8:40 am PST
Subject: Hellooooo???

Hopefully, this will be a successful test message to all of the net using the new Monroe host. If you get this, my return address is cunningb@monroe-emh1.army.mil. If you don't get this, continue to use previous.

Gary Cziko--test message to you earlier this morning indicated host uiuc.edu not registered here, but the full expansion for csg-l may be. Hence dual address. We're getting there.

Bill C.

Date: Thu Mar 26, 1992 8:54 am PST
Subject: Response from Rick on theology

[Martin Taylor 920326]

I posted my comment on Rick's "rant" to him personally, as well as on the net (actually, it was a human-factors problem with the "reply" command). He has replied to me personally, with a permission to post to the net. So here it is.

=========================
Martin

Apparently you posted this to me personally so I'll answer personally. But we could but it on the Net if you like.

You say

>Methinks thou dost protest too much, mine Rick.

Yes, I know. So do others. I get like that. I guess it comes from too many rejection letters based on the same old bullshit -- which is just a rehash of this open loop stuff.

>I have not observed any writer on this group who seemed to me to deny the
>primacy of closed loop effects in the interpretation of behaviour. The
>only question is whether ANY action is possibly open-loop at SOME level of

>the hierarchy.

My impression was that people were trying to show examples of open loop CONTROL. I consider this a factual mistake -- an oxymoron. It get's up my ire because I have had sooooo much experience with reviewers and whatnot dismissing control theory because they KNOW that the important models of control are open loop. Motor programs, dynamical systems, reinforcement theory, etc -- all have been presented to me as the correct explanation of control phenomena. When I see people saying that there are examples of open loop control I suspect that there is a reason that goes beyond scientific analysis.

I don't think the argument was about whether there are ANY open loop connections at some part of the hierarchy. If this were the case the discussion could have ended quickly because any control theorist would agree that at the lowest part of the hierarchy (where the efferents connect to muscles -- and on out -- it is ALL open loop. Muscles cause forces which cause movements which patterns, etc. All cause effect. In fact, open loop causation is what physics, chemistry, astronomy, etc are all about -- and they already have the correct models of these phenomena. I don't think people were trying to say that there is really such a thing as gravity that accelreates objects in a vacuum towarsd earth at 32 ft/s^2. They were saying that causal processes like this could produce controlled results. This is just FACTUALLY wrong (not religious heresy).

> Your rant is way, way, off the point.

I have to believe it was right on target, as usual.

> You take any
>discussion of the possibility that some consistently observed relation between
>a definable stimulus and a describable constellation of actions might be
>open-loop as an assault on your faith. Heretics must be burned!

No way. I NEVER contested the fact that consistent results could be caused by a consistent stimulus. I have no doubt that you could set up a gun in a fixed holder, aimed at the right place and fired it at the bull's eye over and over again WITHOUT LOOKING -- just by pulling the trigger. What I said was that this could not happen in a real environment, where other variables influence the results besides the stimulus (finger on the trigger) and where the connection between stimulus and end result does not remain perfectly calibrated (there is variance in the effect of the hammer, blast, barrel, etc).

Consistent results achieved under the latter circumstance is control -- and it can't be achieved by any open loop system. This is a demonstrable fact; no religion required. Is it heretical to believe that the earth is FLAT if you understand Newtonian physics?. No -- just WRONG. Just as it is WRONG to believe that there is open loop control once you understand control theory.

>It's pretty obvious that you and Bill disagree as to whether there exist
>ANY examples of open-loop constellations of action. The converted are
>usually the most rigid of missionaries.

Nope -- we agree. Open loop components; not open loop control. In the gun example, the bullet's path is not controlled but the ultimate position

of the bullet is controlled (by adjusting where you aim based on previous results). Control is not as good as it could be were the ballistic component also under control.
But the controlled result has far less variance than it would have if it were generated open loop (ie -- if you could not see the result of each shot). Where the bullet ends up is under CLOSED LOOP CONTROL -- the path of the bullet from firing pin to target is determined open loop.

There are tons of open loop components to any control loop -- open loop just means cause effect - THERE IS NO LOOP. In fact, it is the existence of the many "open loop" components of the chain that leads from the outputs of the actor (efferent neural impulses) to the proximal results of these actions (bullet in target) that is one reason why control REQUIRES a closed loop.

I agree that it was probably a tad rude to ask why people are reading the newsgroup -- but, gee, if people don't get this basic point then they are already convinced that the conventional approach is basically OK. The whole point of the control model is that negative feedback is control -- and behavior is control. Open loop explanations are not only unnecessary, they are misleading and factually wrong.

I have been accused of religious zeal about this before. But as I said in other posts, there is really no way to get PCT right and not be seen as a zealot. Unfortunately, PCT shows that convention-al approach is wrong -- ALL WRONG. Bill tones it down pretty well but ask him if my statement is not accurate. If PCT is right, the whole of experimental psychology is wrong. Even if you don't say it, somebody will eventually catch on. But my "rant" had to do with factual errors rather than religious dogma. Open loop systems do not control. That is a fact. They don't sort of control or control sometimes -- they don't control (any more than gravity pulls harder in a vacuum). Wrong is wrong. I don't want to excommunicate anyone -- shit, I despise religion. But I'm afraid I do get a bit impatient with people who keep getting PCT factually wrong after Bill and others have been presenting the facts very clearly for years. I suspect that people have the desire to keep some of the bathwater -- having mistaken it for baby. See Gary Cziko's latest post for a similar speculation -- and Gary is an honorable man, not a fanatic like me.

>Rick, do you want to issue another exclusionary Bull against statisticians?
>Here's your chance.

I like statistics. The argument against statistics is only against using group data as a basis for modeling individual processes. Otherwise statistics is great. I was even planning to do a study of the ability to control statistical variables. Hey, I like to gamble a little too (well, actually, it's not my addiction but I can see the interest).

Feel free to post this reply (along with your original post) to the net if you like. If not, that's fine too.

Best regards   Rick

Date: Thu Mar 26, 1992 8:59 am PST
Subject: Work on CSGnet

[From Gary Cziko 920325.2210]

>(Rick Marken 920325 13:00)

>>Gary Cziko -- what is the CSGNet policy regarding actually doing work
>>over the net? I guess it should be OK -- but it seems like it could
>>actually make work fun. And you know me; I always think something must
>>be wrong if I'm having a good time.

[Martin Taylor 920325 21:00]
>
>Gary, I hope that you would answer "Work on the net is OK, if it illuminates
>or makes concrete aspects of PCT, but not if it is just for the benefit of
>the dyad or small group concerned." As an interested party, I'd like to
>pursue Rick's idea of looking at real concrete HCI problems, but I can do
>that with him off-line if the questions are of insufficiently general
>interest.

Just because I have university post, that doesn't mean that I'm TOTALLY
opposed to work. Go at it publicly and others will probably chime in. If
some others don't like it, they all have delete keys (someone once
said).--Gary

Date: Thu Mar 26, 1992 9:03 am PST
Subject: Mr. Stokes and neural currents

[From: Bruce Nevin (Thu 920326 10:45:14)]

This seems to be right up the CSG alley. Many of you here are better
qualified to respond to Mr. Stokes than I.

Bruce bn@bbn.com

>Really-From: STOKES%MOENG.TOWSON.EDU@BINGVMB.cc.binghamton.edu
>Date: Wed, 25 Mar 1992 15:58 EST
>
> My idea deals with monitoring the electrical impulses in the nuerons and using
>them as input for computer functions, such as movement. I was wondering has
>anybody done something on this (I'm sure someone has) and the specifics of the
>experiment(s). I hope I'm not being too demanding. I thank you for my
>acceptance. I am an electrical engineering student. I might need some help.
>
>Author : Randolph E. Stokes (Is this sign off right?)

Date: Thu Mar 26, 1992 9:06 am PST
Subject: Hellooooo???

Hello yourself???

Date: Thu Mar 26, 1992 9:11 am PST
Subject: Moths in open-loop situation

[From Bill Powers (920326.0900)]

Avery Andrews (920326) --

>But shouldn't we distinguish cases where the action (output = O)'s
>effect on the input (I) flows thru to effect further changes in O from
>those where it doesn't? E.g., in the infamous plummeting moth, it is
>perhaps the case that the change in position of the moth caused by
>folding up (O) causes some change to I, but there will not typically be
>any resulting further change to O ...

I assume that control systems always pass their inputs continuously through the comparator to their outputs as long as they're turned on at all. My picture of the moth's plummeting response is that the moth normally controls for bat-sound intensity by flying away from it (so the inverse-square law reduces the intensity below the reference level -- this would be a one-way system that controls only for excesses of the input over the reference level and does nothing about smaller amounts of input). "Away" might also mean "down," and the control system might have an evolutionary bias for "down" because of the protection afforded by leaf litter etc.

If my concept of the moth's hookup is right, then the plummeting response results when the bat is near enough to cause a very loud sonar sound relative to the reference level, creating a very large error signal. The plummeting response represents the fastest downward velocity that the moth can produce. If the sound is loud enough, however, this response won't be enough to bring the sound intensity below the reference level. The error signal will remain large and the sound input will be uncontrolled because the "output" opposing it can't get any larger. So the moth is in an open-loop situation, even though it is still organized as a control system. It is simply faced with a disturbance larger than it can handle.

This imaginary picture supposes that there are levels of bat-sound intensity below which the plummeting response won't be seen; the moth may descend to the ground or fly away from the sound, but this will suffice to keep the sound below the reference level. This wouldn't be a very dramatic response, because the moth would still be controlling for all the other inputs with which it's concerned and all you'd see would be a bias in the flying patterns away from the sound. With small excesses of sound intensity, you might see the moth descend normally to the ground for a while, then rise again and go about its business. There wouldn't be any big interesting "response" sticking out to draw attention to itself. But you'd be seeing the same control system working in its normal range of operation. This is what I mean by saying that if you only notice extremes of behavior, you'll miss most of what's going on.

>I have no problem at all with the belief that (sub-)systems that are
>supposed to run open loop (my sense) are pretty rare, but I think it's
>quite important to accept them without a fuss if they stand up to
>careful scrutiny. If you don't, people are likely to get the idea that
>PCT is some kind of religion rather than an actual insight into what is
>usually going on with living things.

That's what we're trying to do, subject the behavior to careful scrutiny. Unfortunately we have to imagine a lot of the data, but that could be remedied if someone were to give the moths a closer scrutiny. Maybe they have -- if so I have yet to hear about it. My suspicion is that observers of these moths don't believe that the moth can hear a distant bat sound, and flee from it, unless the moth plummets.

I completely agree with not treating CT as a religion -- forcing the appearance of control onto every situation and rejecting every piece of evidence to the contrary. No matter how much I disbelieve in open-loop behavior in organisms, it can still happen, and I wouldn't automatically reject evidence that says "we looked for control and there wasn't any." If I'm skeptical, it's because I never see data that includes a check to see if control was going on.

Best Bill P.

Date: Thu Mar 26, 1992 10:20 am PST
From: Dag Forssell / MCI ID: 474-2580

TO: csg (Ems)
EMS: INTERNET / MCI ID: 376-5414
MBX: CSG-L@VMD.CSO.UIUC.EDU
Subject: Moths, Smithsonian.
Message-Id: 03920326182030/0004742580NA1EM

From Dag Forssell [920326]

Just recieved the April issue of Smithsonian magazine. An article on "An ancient arms race shows no sign of letting up." includes the following text on moths and bats. A picture is included, which shows a multiple exposure sequence of a moth "dropping" at the approach of a bat. The moth appears not to fold its wings as the armchairs have been presuming, but appears to be actively flying down.

Much here to suggest some rather sophisticated capabilities in these small, simplistic mostly S-R critters. (Just joking as my Swedish friends say).

Quote: ....

Moths and their predators are in an arms race that started millions of years before the Wright brothers made the Dresden raids possible. Butterflies exploit the day, but their "sisters" the moths dominate the insects' share of the night skies. Few vertebrates conquered night flying. Only a small fraction of bird species, mostly owls and goatsuckers, made the transition. Bats, of course, made it their realm. Many species of bats are skilled "moth-ers": they pursue them at speed after detecting them with their highly attuned echolocation system. Some moths, however, have developed "ears" capable of detecting the bat's ultrasonic cries. When they hear a bat coming, the moths take evasive action, including dropping below the bat's track. The parallels of the response of Allied bombers to the radar used by the Germans in World War II are interesting. If we visualize the bombers as the moths, and radars on the ground and in the night-fighter aircraft as bats (a reversal of sizes), the situation is similar. Bombers used rearward-listening radar to detect enemy night fighters. When they detected a fighter, they took evasive action. But heavy bombers, heavily laden, were not very maneuverable. They couldn't dodge about quite as well as moths. Some pilots tried to drop their aircraft into a precipitous dive. Moths also do this; it is easy for them to fold their wings and drop. The next stage in the night-battle escalation is predictable. The night fighter's radar was eventually tuned to detect the bomber's fighter- detector, and thus the bomber itself. Bats have not yet tuned in on moths' ears.

Bombers also used technological disruption. Night fighters came to be guided to bombers by long-distance radars on the ground. The fighters started winning. But nothing remains static. The ground radars could be jammed by various kinds of radio noise. The technological battle swung the other way. Then the fighters acquired radar. Much like a bat, a fighter emitted and listened to radar signals of its own. These, too, proved to be susceptible to countermeasures, however. The RAF could jam the fighters' radar or "clutter" it with strips of aluminum foil. Each bomber in a formation dropped one thousand-strip bundle per mi nute, so that huge clouds of foil foiled the radar. Amazingly, there may be a similar counterweapon among moths. Some moths can produce ultrasonic sounds that fall within the bats' audio frequency. The moths' voice boxes are paired, one on each side of the thorax; double voices must be particularly confusing. Alien sounds in their waveband could confound the bats, exactly in the same way the foil confounded the fighters.

The next steps in the bat-versus-moth war may simply be awaiting discovery by some bright researcher; after all, we did not know a lot about echolocation in bats until after World War Il. My guess would be that the detector will get more complex to meet t he defenses. This may already have happened; bats specializing in moths with ears may have moved to a higher frequency sound outside the moths' hearing range!

As a dedicated spider fanatic, I find the battle between moths and spiders even more fascinating. Web building spiders possess incredibly sensitive vibration receptors. Working a web by night is just the same as working one by day. Eyes are not involved. Moths fill the night skies with potential spider food. They fly through forest gaps that spiders can "fish" with their webs, and they are loaded with the proteins, fats and carbohydrates spiders need for reproduction. They cannot see to avoid webs in the dark.

Moths resemble butterflies in having wings that are proportionately large in relation to body length. This great surface for adhesion should make moths eminently trappable. Moths, however, can escape from the glue on spiderwebs. Their huge wings are covered with loose scales; magnified, these look like tiles on a roof. The scales detach rather easily from most moth wings. When a moth blunders into or even brushes against sticky spider silk, the scales pull off from the wing and stick to the web. The moth may slip free.

Dag Forssell 23903 Via Flamenco Valencia, Ca
91355-2808 Phone (805) 254-1195 Fax (805)
254-7956 Internet: 0004742580@MCIMAIL.COM

Date: Thu Mar 26, 1992 10:14 am PST
Subject: Meditations on Open Loops

[From Rick Marken (920326 9:00)]

I see this morning that there is a long post about open loop behavior from Bill Powers. I haven't read it yet. I decided to post something I wrote last night and make a fool of myself before reading what Bill has to say:

After finishing a personal reply to Martin Taylor I started to ruminate about open loop control (again). The ruminations began with an obvious observation (that I made in my note to Martin). Open loop means NO LOOP. There is a problem with the expression OPEN LOOP: it implies that LOOPS are common and sometimes they are OPEN. I don't like that. I think what is common (in science and in the universe) is cause-effect. That is, NO LOOP. This is what the term OPEN LOOP really refers to; situations where one or more variables have an effect on one or more other variables etc. where none of the variables have an effect on themselves (via any of the other variables). What is less common is where a variable is the start of a chain of cause effect where the last effect in the chain is the cause of variations in the variable in the first part of the chain -- a causal LOOP. When this happens (and in the universe it is apparently rare) the circle of causes and effects has very low gain; cause1 = f(cause1) where f makes the effect of cause 1 on itself quite small. What was amazing was the development of LOOPS that not only have VERY HIGH GAIN but also have NEGATIVE GAIN. So far the only loops of this kind that we know about exist on earth. The search for intelligent life in the universe is a serch for the existence of other high gain, negative feedback loops.

So OPEN LOOP behavior is really a verbal trick. It is an attempt to give a LIFE SCIENCES- like name to processes that have been studied in the natural sciences for decades. The idea that OPEN LOOP processes can be responsible for the behavior of living systems (a behavior we refer to as CONTROL) is exactly the same as saying that causal variables (like gravitational force) can be responsible for the "purposeful behavior" of a leaf as it wafts to the ground. Aristotle was laughed out of the gym for claiming that leafs "seek their natural place" on the ground. But now, 2500 yrs later, life scientists are being celebrated for explaining purposeful behavior with cause - effect models -- ie. physics models. Go figure.

William T. Powers noticed a "fact of life" that should, perhaps, have been obvious to others, given the maturity of control theory as a discipline at the time of his observation. But no one else in the life sciences noticed (or was willing to notice) this fact: when there is high gain, negative feedback from the output to the input of a system, then cause-effect models are no longer appropriate. Powers also showed why negative feedback systems would look like cause-effect systems -- the kind that people assumed they were all along.

There was a BIG PROBLEM with this apparently simple discovery. It was made in the 1960s, about the time that the life sciences in general, and the behavioral sciences in particular, were settled comfortably into a life of studying cause-effect relationships using a statistical/experimental paradigm bequeathed to psychology by R. A. Fisher. What Powers found was that negative feedback made this entire approach irrelevant; it was, quite frankly, ALL WRONG. Who would believe this? It turned out that nobody would (except for a few degenerates from the midwest--like me). Some people did like Powers' language -- hearing the strains of a new approach to that good ol' cybernetics. But many (most?) of those who liked Powers' control theory (other than the crazy degenerates) kept thinking that it cound not really mean what it meant -- that the whole cause-effect kit and kaboodle of psychology had to go. That was too rough. There had to be something good left in all that 100 plus years of work by so many smart, ambitious people. So they tried to preserve what they could -- and, of course, that meant that they could not get the point of control theory. The prime example of this is the Carver/Scheier approach to control theory.

But the fact of the matter is that when there is negative feedback involved in the relationship between an organism and its environment then cause effect laws are no longer applicable, They are not partically applicable, sometimes applicable or sort of applicable. They are WRONG. EVERYTHING that psychology thought it knew would have to GO because it was all based on the assumption that behavior was the last step in a cause effect process. The fact that outputs effect inputs WAS noticed by classical psychology but its existence was thought of as a minor nit. It was just an extra relationship to be considered -- sure, said the psychologists, we know that "feedback" is important. But a lot of behavior is OPEN LOOP, they said -- thus inventing a term that was to haunt control theorists to this day. It suggested that psychologists knew there were LOOPS in behavior -- but it said the existence of these LOOPS was not really THAT important, and, besides, the LOOPS are usually OPEN anyway.

The fact is, there is no other way to go. I don't want to be a radical or "religious zealot" as it appears I am perceived to be. But the uncomfortable reality is that it just can't work both ways -- if organisms are negative feedback control systems, controlling perceptual variables relative to internally specified reference signals, then all the "facts" in your intro psych book are not facts at all. They are, well, illusions that result from looking at a control systems as cause effect systems. They are illusions in the same sense that it is an illusion to look at a falling leaf or rising steam as examples of purposeful behavior.

Sorry, but that's the way it is.

Regards   Rick

Date: Thu Mar 26, 1992 10:29 am PST
Subject: Re: Moths in open-loop situation

[Martin Taylor 920326 12:00]
(Bill Powers 920326.0900)

>If my concept of the moth's hookup is right, then the plummeting
>response results when the bat is near enough to cause a very loud sonar
>sound relative to the reference level, creating a very large error
>signal. The plummeting response represents the fastest downward velocity
>that the moth can produce. If the sound is loud enough, however, this
>response won't be enough to bring the sound intensity below the
>reference level. The error signal will remain large and the sound input
>will be uncontrolled because the "output" opposing it can't get any
>larger. So the moth is in an open-loop situation, even though it is
>still organized as a control system. It is simply faced with a
>disturbance larger than it can handle.

While agreeing with this, it occurs to me that there could be something else going on as well. Unknown to the moth, there may be a communications issue. From the bat's point of view, the beating of the moth's wings provides a particular modulation on the sonar echo that is not provided by a falling leaf. The "evolutionary purpose" of the moth's plummet may not be to get the moth out of the bat's way, because bats are pretty manoeuvrable beasties. Rather, it may be to turn the moth into an inedible leaf. The bat's percept changes from "edible thing in view" to "inedible thing I should avoid."

If this hypothesis is right, then we should be able to model the bat-moth interaction in the same framework as language, as an almost degenerate case. Such cases can be interesting, because they might be tractable without being so trivial as to be worthless.

Martin

Date: Thu Mar 26, 1992 10:32 am PST
Subject: Re: Meditations on Open Loops

[Martin Taylor 920326 12:15]
(Rick Marken 920326 9:00)

I think Rick is right about the connotations of the term "Open Loop," but less correct in asserting that loops are rare in nature. I think it would be extremely rare to find a case in which an event did not have some later consequence for the participating entities. Also, I think it is extremely rare for there to be an identifiable "cause" for an event. Always there is a set of influences, and just as Bill keeps pointing out that desired effects CAN occur without control IFF the environment is very stable, so can a cause be identified IFF the other influences in the environment remain unchanged. The same kinds of arguments and counter-arguments can apply to cause-effect analyses as to closed-loop analyses.

What we probably agree on is that high-gain persistent negative feedback loops probably exist only within living things (or things designed by living things). Unstable ones exist all over the place--just look at the vortices in a turbulent stream. They are negative feedback systems, but they are not control systems because they have no references. The effects of water molecules on one another feed back negatively to maintain a pattern. It's called a "self-organizing system," which is quite different from a control system, though a self-organizing system could (and, I think, did) develop within it sub-parts that are control systems.

Martin

Date: Thu Mar 26, 1992 10:32 am PST
Subject: Some revisons for the Rubber Band Experiment

%%%%%%%%%%%% FROM CHUCK TUCKER 920326 %%%%%%%%%%%%%

This may be "out of sink" but working with my class this semester I found several problems with the RB experiment that uses two coins (one as a target for S and the other a target for E) with the idea that control is indirect in the sense that another will not do something you request unless it serves his or her purpose. The problem is that the other coin (for S's finger) is a disturbance and interfers with the S concentrating on his or her purpose of keeping the know over the coin. My solution is to eliminate the coins and use a diagram with a target with a variety of letters (Clark's suggestion) on the paper where E will attempt to place S's finger. In all trials the S is not aware of their finger over a letter on the paper. This can also be done on the blackboard, sheet of paper on the wall, a transperancy with overhead projector. It is still very portable. Try it. Chuck

Date: Thu Mar 26, 1992 1:23 pm PST
Subject: Diagram for revised RB experiment

FROM CHUCK TUCKER 920326B

I THOUGHT THAT I WOULD TRY TO "DRAW" A DIAGRAM THAT I MENTIONED IN MY PREVIOUS POST"A

```
                RR
                             __+__             K    X
            A       U       |     |
                            +  *  +           O
               R            |__+__|       M           L
        S           Y                           E

                                                    V
                 T                          M
```

HAVE S HOLD KNOW OVER MIDDLE + AND YOU PUT THEIR FINGER IN RR, R, AND T
IN THAT SEQENCE; THEN CHANGE ROLE AND SEE IF S CAN PUT YOUUR FINGER
IN K, X, AND O IN SEQUENCE. TRY IT. CHUCK MUST GO NOW

Date: Thu Mar 26, 1992 11:01 am PST
Subject: Open loop vs. closed loop

[From Rick Marken (920326 9:30)]

Great. Bill's post clarified mine (as usual). Yes. What I was talking about was how closed and open loop situations make a difference in how you deal with the behavior of the same organism. Of course, the situation is typically only closed loop. But there are circumstances (lost steering mechaism, overpowering disturbance) that change them to open loop SITUATIONS -- same organism, new equations.

Bill (920325) says:

>So open-loop systems can exist. Even control systems, in an open-loop
>situation, will behave like open-loop systems: inaccurately and unreliably,

My point precisely. I made it because I want to be sure its clear that adding negative feedback into the relationship between the organism and its environment (into the situation) changes things COMPLETELY. It is not just an interesting detail; its the whole enchilada.

>If any CSGer wants to have the kind of nauseating experience that Rick
>likes to celebrate, read the opening main article in Science for 20 March,
>1992.

No thanks, I haven't eaten yet.   Regards   Rick

Date: Thu Mar 26, 1992 11:49 am PST
Subject: Re: Open- vs. Closed Loop

From Greg Williams (920326)

Rick, in your reply to Martin which he posted, your logic is faultless. No doubt, if someone thinks that "open-loop action is control," then they are SIMPLY WRONG, PROVIDED that they accept your DEFINITIONS that "closed-loop action (only) equals behavior" and "behavior equals control," which reduce to "closed-loop action (only) equals control."

But science is supposed to be more than a matter of definitions. I thought the argument was about whether some organisms use open-loop, pre-calibrated behavioral organizations to achieve what look to observers as adaptations. Contra Bill's latest post (I think), I think the argument was about how organisms are organized INTERNALLY. For example, does Horst Mittelstaedt's preying mantis pre-calibrate its trajectory before lunging and then follow that "aiming" during the lunge, or does it adjust its trajectory as it lunges on the basis of ongoing perceptual inputs? I'm a bit lost by the newest turns (?) in the debate, seemingly toward tautological exclusion of some positions and/or bringing in EXTERNAL (environmental) organization. Maybe I'm just slow to catch on, but I'm still wondering why a person couldn't, when reaching, use BOTH pre-calibration (to get started) AND end-point-perceptual-control (near the goal). After all, when the lights go out, pre-calibration begins to seem pretty important!

Greg

Date: Thu Mar 26, 1992 11:52 am PST
Subject: "Chain" better than "open-loop"?

From Greg Williams (920326-2)

I forgot to mention in my last post that Mittelstaedt, at one of the Macy conferences ("Group Processes," not "Cybernetics") showed three alternative behavioral organizations: his "chain" is like what some netters have been calling "open-loop," his "loop" corresponds to our "closed-loop," and his "mesh" I don't think we need to be concerned with, at least at this stage of the discussion. So there is an early use in biology of a reasonable (I think) substitute for "open-loop" which perhaps avoids some objections of PCTers.

Greg

Date: Thu Mar 26, 1992 12:31 pm PST
Subject: Re: Meditations on Open Loops

[From Rick Marken (920326 11:20)]

Martin Taylor 920326 12:15 says:

>I think Rick is right about the connotations of the term "Open Loop," but
>less correct in asserting that loops are rare in nature. I think it would
>be extremely rare to find a case in which an event did not have some later
>consequence for the participating entities. Also, I think it is extremely
>rare for there to be an identifiable "cause" for an event.

I agree completely and I admit I was wrong -- its mainly very low gain (fractional even) connections between continuously interacting variables.

Now, would a religious fanatic admit he was wrong?

I was also kidding about having fun while working -- I really like it; no guilt at all.

Best reagrds   Rick

Date: Thu Mar 26, 1992 1:57 pm PST
Subject: Re: Open- vs. Closed Loop

[From Rick Marken (920326 12:00)

Greg Williams (920326) says:

>Contra Bill's latest post (I think), I think the argument was about how
>organisms are organized INTERNALLY. For example, does Horst Mittelstaedt's
>preying mantis pre-calibrate its trajectory before lunging and then follow
>that "aiming" during the lunge, or does it adjust its trajectory as it lunges
>on the basis of ongoing perceptual inputs?

What this discussion has made me realize is that the argument is really NOT about how organism's are organized internally. And I think that this is an important point -- and a new realization for me. The negative feedback situation defines the organism as a control system, regarless of how the organism is internally organized. The internal organization of the organism just has to transform sensory inputs into outputs that effect the input in a negative feedback sense. There are probably many organizations that will do this -- and one of these has closed loop systems manipulating the reference for the inputs controlled by other closed loop systems -- that's the HPCT model of the organism, but it is not the only one (as Bill would readily admit).

There is nothing tautological about the fact that negative feedback means that the organism is a control system. The existence of negative feedback connections can be determined by inspection and test. And the fact that the system is controlling input variables can also be determined by inspection. There is no tautology because the inspection could reveal that there is no negative feedback (or controlled input). If there is negative feeback (because of the closed loop SITUATION) then the behavior of the organism cannot be correctly described by an equation of the form o = f(i). It must be described by two simultaneous equations which, when solved for output (o), show that o = -g(d) -- outputs depends on disturbance, NOT INPUT. That is not a tautology, it is a fact. And it is true of ANY organism (no matter what its internal organization) when it is in a high gain, negative feedback relationship with its environment. I am arguing that, if you don't know that this is the case when you are dealing with an organism in this situation (negative feedback) then you cannot possibly come to the correct conclusion about what its internal functional organization might be based on observation of its behavior.

Regards   Rick

Date: Thu Mar 26, 1992 3:21 pm PST
Subject: Re: Open loop vs. closed loop

[Martin Taylor 920326 16:40]
(Bill Powers (I think; I haven't seen the original) quoted by
Rick Marken (920326 9:30))

>If any CSGer wants to have the kind of nauseating experience that Rick
>likes to celebrate, read the opening main article in Science for 20 March,

>1992.

I was going to point out that article, for the same reason. But also have a look at the following one. I think it has a close relation with reorganization.

One thing that might be gleaned from the nauseating article is the vectorial addition of some, but not all, of the neurons involved in the reaching movements. It isn't probably very important, but I think such distributed systems are more likely to be doing real control than are the individuated boxes we draw and discuss here. We do it for convenience, but sometimes people talk as if isolated ECSs actually could be segregated in a living organism. Clean separations may sometimes be possible, but I doubt they are common in higher organisms such as vertebrates.

Martin

Date: Thu Mar 26, 1992 4:41 pm PST
Subject: Re: Open loop vs. closed loop

[Martin Taylor 920326 19:00]
(Bill Powers 920325 20:00)

I'm afraid I'm an addict of CSG-L...help!!

>When an organism that normally acts as a control system experiences an
>open-loop situation, it becomes hyper-responsive to changes in its input
>because those changes are no longer counteracted; they show up directly as
>error signals. There will be hyperresponsivity to changes in reference
>signals, too, because the reference saôoal also shoºÄ\]+ directly as a
>change in error signal. Normally the perceptual signal would immediately
>catch up and the error would be kept from getting large. But with the loop
>open, the perceptual signal no longer changes because of the output, and
>the error signal remains large.
>

If you remember, I posted a note from the paper about the results of the last space shuttle flight. I think it worth re-posting in light of the above. As yourself whether, as a CSGer, you would use the fourth word of the quote:

"One early and surprising result of Dr. Watt's experiments was that it appears muscular changes that on Earth keep the human body in a stable, upright position increase rather than decrease in weightlessness."

Martin

Date: Thu Mar 26, 1992 4:49 pm PST
From: Dag Forssell / MCI ID: 474-2580

TO: csg (Ems)
EMS: INTERNET / MCI ID: 376-5414
MBX: CSG-L@VMD.CSO.UIUC.EDU
Subject: correction
Message-Id: 14920327004941/0004742580NA2EM

I was typing on line to MCI mail and a part got left off in reformatting / sending. Here is corrected retransmission.

From Dag Forssell [920326 - 2]

Gary asked about "Tightly coupled". Imagine that you are playing the rubber band demonstration with a strong machine, programmed to go through a set pattern of motion.

You have no difficulty, since the machine only influences the position of the knot, through the tension in the rubber band.

Now "tightly couple" the knot to the machine by substituting a stick (or a rope as I like to do when demonstrating conflict between two "pullers"). If you still are connected to the knot by a rubber band on your side, you will pull in vain.

If the stick is extended to your hand, you will be pulled along, powerless to do otherwise. You are being controlled by the machine with overwhelming physical force, the only way Bill says you can be controlled.

I believe it is important to remember one of the hallmarks of control systems: amplification. This term does not communicate well. I am shifting my language to "the direction of resources" or something like that, with emphasis on resources. Your heating system at home opens a valve to release (and ignite) a stream of natural gas. The stream is not finely calibrated, but has a powerful influence on the air temperature.

If you have an air conditioner working at the same time, set at 68 degrees, while the heater is set at 75, the two will pull (with tight coupling) on the air temperature knot with as much influence as each is capable of.

If the gas line has the capability to release more resources to raise the temp than the air cond has resources to lower it, then the air temperature will stay at 75 degrees.

The rubber band is such a marvelous tool, because it shows influence without tight coupling. Try the demonstration with a rope, and two dots. One towards the left as a target for the left person and one a little to the right (one foot apart if you are at a blackboard with a 4 foot rope, which works best, one inch if you are on a paper with a short string) as a target for the person on the right. See which person is willing tu pull hardest. This person will pull the knot to his dot and keep it there. This will illustrate the heater / air cond conflict.

A detail: The heater and the air cond are separate control systems, as in wall heater and window air cond, with separate thermostats.

Hope this helps. I enjoy your questions.

Dag

Date: Thu Mar 26, 1992 4:58 pm PST
Subject: Open loop; dropouts; consistent results

[From Bill Powers (920326.1600)] --

Open Loop:

The term "open loop" is another of those control-engineer's terms that has been picked up (like "feedforward") by non-technical types and interpreted by free association or verbal logic. The original usage of "open loop" was in the context of measuring the characteristics of control systems. In a closed configuration with high loop gain, there is no way to measure the loop gain directly, because at any point in the circuit, only one signal amplitude can be measured at one time. Intuitively, if you follow an imaginary perturbation all the way around the loop to its starting point, you imagine a perturbation of much larger amplitude and opposite sign appearing at the end of this complete trip. If course this doesn't happen -- unless you physically cut one of the signal paths, inject a signal into the downstream broken end of the wire, and measure what comes back at the upstream broken end. When you actually break the loop, you can see the loop gain and reversal of sign -- but of course the system stops working.

So the concept of "open loop" really has meaning only when the loop is normally closed, but is deliberately broken in order to measure the loop gain directly. A system that isn't designed to run with the loop closed isn't an "open-loop" system -- it's just an input-output device.

Martin Taylor (920326.1200) --

>What we probably agree on is that high-gain persistent negative >feedback loops probably exist only within living things (or things >designed by living things). Unstable ones exist all over the place-->just look at the vortices in a turbulent stream. They are negative >feedback systems, but they are not control systems because they have no >references.

There's another factor beside the reference, which is that "high gain" of which you speak. Systems like vortices or pendulums do not have high loop gain: they are dissipative systems in which energy decreases (or entropy increases) continuously as you trace the path around the loop. There is no process in this loop that entails a large power GAIN -- where the rate at which energy comes out is far greater than the rate it comes in. The power gain needed to sustain natural oscillations is usually just a little greater than unity.

In a control system there is always at least one element of the loop that provides for a very large power gain. This is usually the effector of the system, which draws on metabolic energy supplies as directed by the low-energy signal that operates the effector. In neural systems there is power gain everywhere that neural signals are generated and transmitted; in sensors it can be quite large, but even in signal transmission it is enough to make up for the energy losses entailed by pumping ions around. In general, though, the power gain in the effector -- the muscle -- is by far the largest one involved in a given control loop. I would guess that the gain is at least in the millions going from motor nerve output signals to output work done by a muscle. I don't think that this amount of continuous power gain is found in any natural system but a living organism (on the same scale of size).

you say

>I think it would be extremely rare to find a case in which an event did
>not have some later consequence for the participating entities.

This doesn't mean that negative feedback control is universal. First, there has to be an important amount of loop gain in the "later consequence." Following the loop all the way around starting with the consequence, the effect of the consequence on the same consequence via the behaving system must be, in comparison with the original perturbation, MUCH LARGER and of OPPOSITE SIGN if negative feedback control is to exist. That's far from common in nonliving systems.

Second, it matters HOW SOON the "later consequence" occurs. If I press an elevator button and one hour later the elevator arrives, the causal connection between the action and its result becomes very uncertain (especially if 40 other people have also pressed the button in the meantime). For control to work, the reflected effect of output on input must be systematic and clearly so. If the effect is very delayed, the sensor must contain smoothing and averaging of corresponding slowness in order for stability to be preserved with high loop gain. This smoothing also tends to average out brief chance effects on the input from other sources: they don't qualify as disturbances on the relevant time scale. In general, the time-scale of the control loop must be adjusted so that the feedback effects occur WHILE THE ACTION IS OCCCURING, with the proviso that the term "while" is modified by the amount of slowing and averaging that goes on in the system. When the reflected effect is long delayed, not only is precise control made difficult, but interference from independent disturbances masks the feedback effects, requiring further filtering and slowing.

In the main types of control systems we see in behavior, the feedback is strong and concurrent with the sensory changes in the system.

Gary Cziko asked whether people who fully understand the CT model ever renounce it and go back to their former beliefs. While Mary's answer is right -- I tend to lose track of these people and don't know their reasons for dropping out -- I have known some people well enough to guess at their reasons. I haven't seen anyone drop out who knows control theory well enough to field random questions about it and come up with the same answers that any other well-versed CSGer would produce.

I have, however, seen people drop out short of reaching that point. Some people drop out because they just don't get the basic ideas -- when they look into control theory they don't find what they thought they would find, and don't care to go any further. Others drop out as soon as they see that control theory doesn't support the beliefs they already hold: they simply aren't willing to try on a different concept. The most tantalizing cases are the people who do go quite far toward getting all the basic ideas, yet have a personal agenda they're not willing to part with. We've seen people like this drop in and out of the net; those who were around know who I mean. There comes a point when anyone entering control theory from some conventional discipline has to realize that if control theory is right, the teachings of his or her discipline are most probably fundamentally wrong. If there's a large investment in the traditional discipline, and good understanding of control theory, the conflict is made all the more extreme. I can't say I blame such people for finally deciding to back off and return to familiar ground. It's a relief to everyone when the pressure is finally off.

But there is definitely a point of no return. Once you grasp what a control system is, and realize that what you are doing is not just "behaving" but controlling, there's no way to forget that.

Last remarks (I'm just about written out for now). I think we may have been overlooking a difference in meanings when different people talk about organisms producing "consistent results." Under traditional approaches, a response is considered to be a consistent effect of a stimulus if there's a "high" correlation between them -- say, 0.8. In any control-system experiment, a scatter plot of input and output points that showed a correlation of only 0.8 would be considered a sign of a busted experiment. You have to see such a scatter plot to realize just how bad data have to be to get such a low correlation. Looking at various individual points, you can find ratios of output/input that vary by hundreds of percent, by factors of 50 or more -- and the hypothesis is that all data points represent the SAME ratio.

I used to have a martin house (blue martins, not taylor martins) in my back yard. Those birds often wouldn't bother to use the perch: they'd fly up to the hole, fold their wings, and pop in without touching the sides. Consider the precision of control that's needed to do that from an infinitely variable set of starting positions and trajectories. How accurately do you think the aerodynamic forces of the wings had to be adjusted to make a twenty foot approach curve on a breezy day go within a quarter of an inch of the center of the hole? I never saw a bird miss a try. A scatter plot of disturbances versus compensatory forces would have been a straight line.

Or look at a hummingbird. It is frozen in the air near the flower. When it's feeding, its head and beak are nailed to the flower in absolute immobility, while the body moves slightly around, and the wings continue their blur of action. How precisely do you think the forces from the beating wings and the little body movements cancel disturbances from the wind? It's hard to see any movement of the head at all; the bird can detect and correct smaller errors than you can see. And then the bird suddenly backs off six inches to a new frozen position, swoops twenty feet and runs into another invisible stone wall, and so on as long as you care to watch it. What are you going to do with this kind of data? Count flowers visited as a function of flower type?

Organisms are constantly producing this sort of precise control of consequences of acting. If the correlations involved were as bad as 0.8 we'd be finding dead bodies all over the place. God knows what the real correlations would be -- strings of nines.

The problem with conventional approaches to behavior is that scientists look at examples of exquisitely precise control going on all around them and don't see anything happening. A conventional study of martins entering their martin house holes would ignore the control process being demonstrated hundreds of times per day, and would search for some relationship between, say, number of hole entries per hour and proximity of human dwellings. Of course by focussing on such randomly selected bullshit relationships, they manage to come up with correlations of 0.8, or usually a lot less. That's because there really isn't any important relationship there -- just some sort of little fringe effect, which can then be blown up into significance by suitably far-fetched hypotheses and grandiose generalizations. The main effect, the striking effect, the effect that reveals volumes about the real nature of martins, is just taken for granted and counted as hole-entries per hour. Miracles per hour would be more like it.

All the important information about the nature of living systems is contained in the details of control behaviors that conventional scientists simply name and tick off as "occurrances." You can take any one of those behaviors -- a rat jumping off its perch after a puff of air and landing on a small shelf while killing its momentum and regaining balance -- and by studying it in detail, find out more in an hour about how that behavior works than you can in a lifetime of counting events and searching for suggestions of possible relationships. The conventional sciences of behavior have just been totally on the wrong track, thinking they were learning something about behavior when in fact they were learning nothing. NOTHING.

I get very impatient when I hear people throwing up "facts" about behavior that have been discovered in the conventional ways. When you track them down, you end up with those scatter plots again, and realize that any individual organism could serve as a counterexample. I don't think that we who are interested in control processes have any obligation to try to match our theory to facts that are probably untrue of most individual organisms, and don't mean anything even when they are true. I am fed up with conventional approaches to behavior. I want to associated with people who are doing something real. And I'm old enough to get away with it.

Best regards,   Bill P.

Date: Thu Mar 26, 1992 5:38 pm PST
Subject: Re: Open loop; dropouts; consistent results

[Martin Taylor 920326 19:30]
(Bill Powers 920326.1600)

I don't want to get into an irrelevant discussion of self-organizing systems, but I do think you want to be a bit careful about the entropy considerations of far-from-equilibrium energy flows. There is the potential for quite strong gains in the feedback loop that maintains the stability of structures in such a flow. It is the flow that permits the dissipative system to maintain its entropy. If you go around the loop, you see that there isn't really very much difference from a true control system except for that variable reference (or rather any reference at all). There's no comparator, just a negative feedback loop that stabilizes a structure, and the structure may be a sequence just as readily as a visible structure.

I don't think the above is an important point in the discussion, but your later discussion of the timing effects is important.

You say "it matters HOW SOON the "later consequence" occurs." Well, yes it does, but not in the way you discuss. As you go to higher levels in the hierarchy, delayed effects and effects strung out in time become very common both as reference and as percept. We do not know exactly what we did to cause what we are now perceiving (and read "cause" with a large amount of salt). We try something else because the percept is not changing satisfactorily, even though God might know we had been doing the right thing all along. You have yourself said that it is unclear how to treat error signals at the sequence level. And things only get worse at higher levels.

At high levels, there are SEVERE statistical problems in discovering whether we are controlling what we think we are. We have to rely on the imagination loop without real-world testing for a lot of our error-correcting behaviour (it's called planning and prediction). If you were correct that "For control to work, the reflected effect of output on input must be systematic and clearly so." then we could hardly talk about high-level control systems at all, because the environmental disturbances are very great and prolonged. But it is also true that control must be effective while the action is occurring, which means a lot more than just filtering and averaging percepts.

You make nice examples, but they are of low-level control. Why shouldn't the ethologist try to determine whether the bird is controlling for the neighbourhood of human habitation? Lorentz (I think) did careful disturbance experiments on the visual neighbourhood of sites where birds nested (?), and found out what they controlled for. But such an experiment need not always work, because in a complex situation the experimenter will not necessarily hit on all the entities that form part of the controlled perceptual variable. The results will yield correlations well below your 0.99, and yet they will be valid and valuable. Any reliable correlation says that the things studied project to some extent on the direction of interest in the multidimensional perceptual space. I really cannot believe that you will maintain the following when dealing at high levels of control: "In any control-system experiment, a scatter plot of input and output points that showed a correlation of only 0.8 would be considered a sign of a busted experiment." That may be valid for muscular tracking experiments, but psychotherapy experiments? Come, now!

I commend once again the article in Science about granular systems.

Martin

Date: Thu Mar 26, 1992 5:39 pm PST
Subject: chaos, different worlds

[From: Jeff Dooley 920326.1500]

(Martin Taylor 920324.1900)

Thanks for reminding me of the Skarda-Freeman article, "How Brains Make Chaos to Make Sense of the World," in _Behavioral and Brain Sciences_ (June 1987). Looking back over the article, I'm struck by how much of it I failed to understand on first reading. In retro, after conceptually adding control dynamics to the mechanism that may be said to "achieve" chaos in neural activity, this "achievement" makes more sense as a product of control. Reference would be set (in individuals), on this conjecture, through the lawlike (control?) activities of the ensemble's generic organizing function (Kauffman's model) manifesting ontogenetically. But why chaos? Actually, here it seems to help to have Kauffman's notion of the "edge of chaos"--a distinction I do not find explicit in Freeman (perhaps I just missed it?) Anyway, Skarda and Freeman spell out the value: *adaptability, or the ability to accommodate and learn.* Quoting from p. 171:

> "Without chaotic behavior the neural system cannot add a new odor [sensation, vector, second level output?] to its repertoire of learned odors. Chaos provides the system with a deterministic "I don't know" state within which new activity patterns can be generated. . .If the odor is novel and the system does not already have a global activity pattern corresponding to the odor, then instead of producing one of its previously learned activity patterns, the system falls into a high-level chaotic state rather than into the basin for the background odor. This "chaotic well" enables the system to avoid all of its previously learned activity patterns and to produce a new one."

Obviously, a cognitive system whose livelihood depends upon the inductive establishment and taxonomy of perceptual vectors (as I think ours does) would benefit from achieving a degree of structural plasticity like that of chaos and then to maintain equilibrium around it. Is this possible "evidence" of control activity on, perhaps, the phylogenetic level of organization (as Kauffman would allow)? The authors even suggest that "chaos is controlled noise with precisely defined properties (p. 165)." But I don't see reference to *how* or by what mechanisms such control may be explained. They invoke local feedback (same-level, as you point out?) as a mechanism by which bulb-wide activation vectors are enabled, but I don't see further discussion of higher-order feedbacks. These are all blank spaces that it seems PCT could help fill in.

(Gary Cziko 920325.1720)

Kuhn could only say, "In a sense that I am unable to explicate further, the proponents of competing paradigms practice their trades in different worlds" (_Structure of Scientific Rev._ p. 150). Your simple notion of these proponents controlling for their different (incommensurable) perceptions potentially rescues the whole concept of "theory-laden" observation from the intentional depths where it has lurked out of empirical reach for nearly 50 years.

Theory change. He remarks on the dynamics of theory change with the question: which programme solves the more significant question?. It seems that "significant" here has contingent, pragmatic overtones, and is, in any case, a problematic predicate. A new theory might also be judged superior for, according to Kuhn, its 1. accuracy, 2. consistency, 3. broadness of scope or generality, 4. simplicity ("bringing order to phenomena"), and 5. fruitfulness, as in explanatory power over previous anomalies and as in disclosure of novel relations or predictions of phenomena. (Essential Tension, p. 322). Beyond these lies the subjective component--value-laden rationality--that plays a crucial role in theory selection, as I take it, regardless of the five criteria above. This notion lands us right back in the explication muddle that Kuhn could only shrug off but which your suggestion seems in a position to help clarify. I'd like to explore this further but parsimony forbids.

If it were possible to find someone who had grasped PCT (how to determine?) and then went back to S-R we could ask "why?" and gather a wealth of information on these questions of theory choice, objective criteria, subjective criteria.

jeff dooley   dooley@well.sf.ca.us

Date: Thu Mar 26, 1992 5:39 pm PST
Subject: Re: Open/Closed Loops

From Greg Williams (920325-3)

>Rick Marken via Martin Taylor 920326

>My impression was that people were trying to show examples of open loop
>CONTROL. I consider this a factual mistake -- an oxymoron.

Only if YOUR definition of control is accepted.

>I don't think the argument was about whether there are ANY open loop
>connections at some part of the hierarchy. If this were the case the
>discussion could have ended quickly because any control theorist
>would agree that at the lowest part of the hierarchy (where the efferents
>connect to muscles -- and on out -- it is ALL open loop.

The question is the extent to which there are open loops ABOVE that level.

>I NEVER contested the fact that consistent results could be caused
>by a consistent stimulus. I have no doubt that you could set up a gun
>in a fixed holder, aimed at the right place and fired it at the
>bull's eye over and over again WITHOUT LOOKING -- just by pulling
>the trigger. What I said was that this could not happen in a real
>environment, where other variables influence the results besides
>the stimulus (finger on the trigger) and where the connection between
>stimulus and end result does not remain perfectly calibrated (there
>is variance in the effect of the hammer, blast, barrel, etc).

Perhaps you are overestimating the desirability of high-precision control in many circumstances? Maybe some organisms can get some things done with sloppy calibration and considerable drift, and it's GOOD ENOUGH for evolutionary stability. It could be that the extra apparatus required by closed-loop organization is less "economical." The belief in closed-loops everywhere strikes me as quite anthropomorphic -- we do value our own high-precision closed-loop control, don't we? The evolved organism's burden... or pathetic fallacy?

>There are tons of open loop components to any control loop -- open
>loop just means cause effect - THERE IS NO LOOP.

Yet, consider this. Suppose there is a calibration loop which tunes up, over many trials (based on each trial's outcome), the open-loop action. There might not be a continuously acting loop in any one trial, but there is a bigger loop resetting the non-loop's pre-calibration. If you count this as closed-loop control OVERALL, then why gripe about folks who are interested in its open-loop component?

>If PCT is right, the whole of experimental psychology is wrong.

How could this be, if open-loops can be parts of closed-loops? Somebody who looks at the open-loops ONLY might be missing the big, important picture, but how does that make their data wrong?

>Rick Marken (920326 12:00)

>What this discussion has made me realize is that the argument is really
>NOT about how organism's are organized internally.

Then let's start a new discussion about how organisms are organized internally -- specifically, about whether certain actions are due to "chains" or to "loops." By "chains," I mean that INTERNALLY there are no perceptual signals going to comparators, resulting in error signals which influence the actions' "trajectories" during those trajectories; generally, I would expect "chains" to look like precalibrated "responses" to trigger-like "stimuli". Such "chains" could be "output functions" of loops, and all I think would be sacrificed in PCT is the idea that control is always continuous.

>The negative feedback situation defines the organism as a control system,
>regarless of how the organism is internally organized.

But the internal organization has to meet certain requirements for the negative feedback situation to exist!

>The internal organization of the organism just has to transform sensory inputs
>into outputs that effect the input in a negative feedback sense.

Some people think that this is a big "just," more suitable for empirical investigation than decree.

>There is nothing tautological about the fact that negative feedback
>means that the organism is a control system.

It is if you add "only" before "negative feedback."

>The existence of negative feedback connections can be determined by inspection
>and test. And the fact that the system is controlling input variables can also
>be determined by inspection. There is no tautology because the inspection
>could reveal that there is no negative feedback (or controlled input). If
>there is negative feeback (because of the closed loop SITUATION) then the
>behavior of the organism cannot be correctly described by an equation of the
>form o = f(i). It must be described by two simultaneous equations which, when
>solved for output (o), show that o = -g(d) -- outputs depends on disturbance,
>NOT INPUT. That is not a tautology, it is a fact.

NOW YOU'RE TALKING. RIGHT ON! But that isn't what you said earlier.

Greg

Date: Fri Mar 27, 1992 8:05 am PST
Subject: bounced reply to Rick

[Martin Taylor 920227 1010]
I sent the following to Rick, having intended it for CSG-L. So I sent it to CSG-L but got the address wrong and it bounced.   But I also sent Rick's response, and it didn't bounce, so you should see what it was he was responding to. So here, belatedly, it is.

Isn't the maintenance of a reference in the face of disturbance wonderful?

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

[Martin Taylor 920325 17:00]
(Rick Marken 920324 18:00)

Methinks thou dost protest too much, mine Rick.

I have not observed any writer on this group who seemed to me to deny the primacy of closed loop effects in the interpretation of behaviour. The only question is whether ANY action is possibly open-loop at SOME level of the hierarchy. I suppose there may be some minor disagreements about whether there exist open loops that are closed by evolution rather than within the living organism. Your rant is way, way, off the point. You take any discussion of the possibility that some consistently observed relation between a definable stimulus and a describable constellation of actions might be open-loop as an assault on your faith. Heretics must be burned!

Speaking for myself, I think I have pretty completely absorbed and integrated into my natural ways of thinking the concept that all behaviour is in the end the control of perception. It makes so much of what we observe people to do intelligible. But there are cases, such as target-shooting, where the control is exercised over an outer loop that contains NECESSARILY open-loop action elements. You can't correct the path of the shot bullet, so you had better predict where it will go before you press the trigger. You can certainly control the aiming point through the sights IF the sights are reliably fixed (massively, as Bill puts it when describing possible conditions for open-loop behaviour to be effective) to the gun. But only the next bullet's strike can be affected by perception of the results of the last shot. Open-loop, predictive actions are a part of this whole game. The fact that all behaviour comes down to the control of perception doesn't falsify the notion that there can be pretty complex open-loop constellations of actions.
>
>Here are some test questions for those of you who think open loop
>processes are important in some way.
>...
>What is your reason for even reading this newsgroup?

Could it be that the great unwashed might want to learn why they might want to be baptized?

It's pretty obvious that you and Bill disagree as to whether there exist ANY examples of open-loop constellations of action. The converted are usually the most rigid of missionaries.

To put another cat in the neighbourhood of a pigeon roost, Bruce Nevin a few days ago said something to the effect that PCT justified his failure to learn about statistics in his youth. Rick and Bill (Bill less so) have both inveighed against statistics. It is my belief that you CANNOT understand PCT, and especially reorganization, without a reasonably intuitive understanding of statistics. Any attempt to develop PCT deeply, beyond simple tracking tasks, is going to founder on improper appreciation of the statistical problems involved in the perceptual functions, if those problems are simply defined out of existence. My modularity argument for reorganization comes directly out of statistics, and I think most of the structural features of the hierarchy will ultimately be derivable from the statistics of the environment and the purposes that organisms can try to achieve.

Rick, do you want to issue another exclusionary Bull against statisticians? Here's your chance.

Martin

Date: Fri Mar 27, 1992 8:48 am PST
Subject: Nonliving Closed Loops; "Motor Control"

[from Gary Cziko 920327.0900]

Recent posts from Rick Marken, Bill Powers, and Martin Taylor (too much trouble to look up all the dates) have led me to the conclusion that closed loops do exist in non-living, non-human-made situations, but that they don't have BOTH negative feedback and high gain. Is this right?

A pendulum has negative feedback but the energy that goes into restoring the pendulum does not exceed the energy that went into disturbing it (I suppose the restoring energy has to be less because of energy loss along the way). Nuclear fission would seem to have high loop gain, but it is positive feedback (the reaction tends to "run away"). But only living systems or engineered systems made to simulate living systems have BOTH negative feedback and high loop gain.

Is this on the right track? Let me know.

==

Bill Powers (920325.2000)]

>When we consider locomotion, the arrival of an organism at
>a particular place, or even the placement of its limbs in a particular
>orientation, results from the application of muscle forces to limbs, and
>the subsequence effects of limb forces on other objects. To go from forces
>to positions requires two time integrations, nonlinear ones when jointed
>limbs are involved. Time integrations are notoriously sensitive --
>hypersensitive -- not only to initial conditions, but to very small errors
>of computation.

Let me see if I, a very non-physicist, can understand this. I suppose you mean that the forces muscles produce are accelerations. So for a limb to wind up in a desired location, the acceleration is integrated by the laws of physics to produce velocity and then the velocity is integrated to give position. And there is lots of slop along the way in that tiny errors in initial accelerations cause large errors in velocity and likewise from velocity to position. This is why when you spin the Wheel of Fortune by applying an acceleration it is very difficult to get it to come up with the number you want if you make turn more than 1/2 revolution or so--the result is essentially random.

Is this what "motor control" researchers actually believe we do--throw our limbs around with pre-computed initial forces and then see where they end up? This seems so obviously absurd I can't believe anybody would hold such a view. At the least there must be SOME checks along the way to see if the limb is getting where it's supposed to be, or perhaps something like Greg Williams proposes, an initial calculated push followed up with feedback control toward the end of the move. If people put forth such purely open-loop models, wouldn't it be very easy to show how such a model can't work due to the laws of physics?--Gary

P.S. I just realized that there are some "sports" which seem to depend on skill in turning accelerations into positions--like shuffleboard and curling. So I suppose a fair degree of accuracy can be developed. But if it weren't hard I suppose some people wouldn't find it interesting.

Gary A. Cziko

Date: Fri Mar 27, 1992 10:02 am PST
Subject: Midwest Sociologial Meeting

[from Gary Cziko 920327.0945]

Clark McPhail recently filled me in on what's happening in Kansas City in April. I thought the rest of the network would like to see this as well.--Gary
===
[from Clark McPhail]

April 1-4 are the dates of the Midwest Sociological Society annual meeting. The theme this year is Society and Individual, a theme set by, guess who, me. The highlight of the meeting as far as I am concerned, since I organized this and will preside, will be the plenary session on Friday afternoon, April 3, at 4:30pm: "Individual and Society: An Alternative Perspective." The first presentation will be by Richard Marken, "A perceptual control theory analysis of the individual in society." The second presentation will be by Thomas Bourbon, "A perceptual control theory analysis of individuals in cooperative behavior." The third presentation will be by Kent McClelland, "Implications of perceptual control theory for a sociological understanding of individual and society." The discussant will be Charles W. Tucker. There will be computer projections of simulations by Marken and Bourbon. We hope people will be intrigued, provoked, and even interested in learning more about PCT.

cheers,   Clark

Date: Fri Mar 27, 1992 10:13 am PST
Subject: MPD

To the clinicians and anyone else,

This is the second time today I've typed this so it has to be brief:

I have a friend(s) with Multiple Personality Disorder. There is a dominant personality (A) and a few other major ones (one of which is a friend of mine--B) and a whole bunch of minor ones. They have distinct personalities, not just different traits of the same person. Different voices, ages, mannerisms, interests, abilities, knowledge, attitudes, goals, etc. The artistic abilites of A are not present in B, and the athletic abilities of one do not transfer to another. One of the personalities is 5 years old--talks, thinks, and feels like a 5 year old.

A blow to one of the temporal lobes may be a partial cause. Dramatic (occultic) abuse in early childhood is certainly involved. Epilepsy and narcolepsy may be involved also. Person B is co-conscious with A, meaning that while A is "out"(interacting with the world), B can observe and think and feel also if she so desires. A is not co-conscious with B--A simply "disappears." I think A is the only one of the major personalities who are not co-conscious. This means that B can do things without A ever knowing.

I have many questions concerning MPD and PCT. First, if person B wants A not to see something (say, three words in a message) she can make A not "see" those words. How might this fit into PCT?

Second, A can make decisions, but since B has her own desires, she can influence the outcome of the decisions without A's awareness of it. (I asked B why A called me the other day and she said that A _thinks_ that she made the decision when in actuality B influenced her to make the decision to call me). How might one model that?

Third, is MPD "simply" a multiplicity of self-concepts residing at the system level? Just as I satisfy my desire to feel achievement by, say, being a good basketball player or writing good papers or cleaning the apartment, is there one thing being fulfilled by these multiple self-concepts or not? Is it OK to equate self-concept with personality? What keeps this from being more prevailant? Is it only random variation and selection probabilities which make most of us have One personality (i.e. it's not necessary, only common that there is one to one correspondance between personality and body). There are advantages to having multiple personalities--escape is certainly one: if you don't like what's happening then disappear and let someone else come in. If someone is mad at A, she can leave and B could appear--that someone cannot be mad at B. But there are certainly disadvantages--holding a job is tough (but I don't think that is a necessary problem). How one wants to decorate the apartment seems to be a real problem.

Well, I could say a lot more, but I want to hear some clinical opinions first, so my opinions can be more informed.

Fortunately, person A loves psychology and finds herself to be a very interesting subject.

Please send a copy of the reply to my personal mailbox (m-olson@uiuc.edu)--I don't always get to check the net.

Mark Olson

Date: Fri Mar 27, 1992 11:00 am PST
Subject: MPD

Gary,
I spent a half hour writing what I'm rewriting below. I typed on the news group program and when I went to post it it said it couldn't be posted--how can I avoid that happening again?!
Why might Bill not be able to send messages to my personal mailbox? Does m-olson@uiuc.edu only work for U of I?

To the clinicians and anyone else,

This is the second time today I've typed this so it has to be brief:
I have a friend(s) with Multiple Personality Disorder. There is a dominant personality (A) and a few other major ones (one of which is a friend of mine--B) and a whole bunch of minor ones. They have distinct personalities, not just different traits of the same person. Different voices, ages, mannerisms, interests, abilities, knowledge, attitudes, goals, etc. The artistic abilites of A are not present in B, and the athletic abilities of one do not transfer to another. One of the personalities is 5 years old--talks, thinks, and feels like a 5 year old.

A blow to one of the temporal lobes may be a partial cause. Dramatic (occultic) abuse in early childhood is certainly involved. Epilepsy and narcolepsy may be involved also. Person B is co-conscious with A, meaning that while A is "out"(interacting with the world), B can observe and think and feel also if she so desires. A is not co-conscious with B--A simply "disappears." I think A is the only one of the major personalities who are not co-conscious. This means that B can do things without A ever knowing.

I have many questions concerning MPD and PCT. First, if person B wants A not to see something (say, three words in a message) she can make A not "see" those words. How might this fit into PCT?

Second, A can make decisions, but since B has her own desires, she can influence the outcome of the decisions without A's awareness of it. (I asked B why A called me the other day and she said that A _thinks_ that she made the decision when in actuality B influenced her to make the decision to call me). How might one model that?

Third, is MPD "simply" a multiplicity of self-concepts residing at the system level? Just as I satisfy my desire to feel achievement by, say, being a good basketball player or writing good papers or cleaning the apartment, is there one thing being fulfilled by these multiple self-concepts or not? Is it OK to equate self-concept with personality? What keeps this from being more prevailant? Is it only random variation and selection probabilities which make most of us have One personality (i.e. it's not necessary, only common that there is one to one correspondance between personality and body). There are advantages to having multiple personalities--escape is certainly one: if you don't like what's happening then disappear and let someone else come in. If someone is mad at A, she can leave and B could appear--that someone cannot be mad at B. But there are certainly disadvantages--holding a job is tough (but I don't think that is a necessary problem). How one wants to decorate the apartment seems to be a real problem.

Well, I could say a lot more, but I want to hear some clinical opinions first, so my opinions can be more informed.

Fortunately, person A loves psychology and finds herself to be a very interesting subject.

Please send a copy of the reply to my personal mailbox (m-olson@uiuc.edu)--I don't always get to check the net.

Mark Olson

Date: Sat Mar 28, 1992 10:43 am PST
Subject: loops, selection

Some of the things being said about loops now are getting to be a bit more like some of the ideas I was trying to push a few months ago (`e-control'), so I'll try another foray in that general direction.

DNA variants do better or worse due to their differing abilities to maintain conditions suitable for their replication, so what is fundamental is what I'll call `condition maintainence'. Usually, the best way to attain this is to have a perceptual function P that approximately perceives the relevant condition C (exact 100% reliable perception tends to be too expensive, if it's even possible), and a feedback loop controlling the value of P, and thereby, indirectly, maintaining C (which will be more or less tightly coupled to P). So control systems typically get selected for on the basis of their ability to maintain conditions, which either contribute directly to DNA replication (Darwinian selection), or to the control of perceptions of superordinate systems (selection by (some kind of) reorganization), which are themselves being selected for on the basis of their ability to maintain conditions.

My picture would be something like this:

```
                  -------------------------
               I |        Critter          |
  --> ---------->|-----                    |
     ^           |     |                   |
     |           |   P v     Ref           |
  -->C           |    ECS<--------         |
     ^           |     |                   |
     |         O |     |                   |
      <----------|<----                    |
                 |                         |
                  -------------------------
```

What gets controlled by the ECS is some function of I, designated by P, but what the system is being selected for doing is maintaining some circumstance C. But due to possibility of disturbances in between C and I, or limitations on available perceptual and computational resources, what actually gets controlled is not always exactly what ought to be (though in systems with a good reorganization facility and plenty of perceptual and computational resources, P will track C *very* closely). Nota Bene: C is not something that the organism is aware of, but is rather the reason (as determined by our analysis of the whole situation) for the survival or otherwise of the organism, or the success or failure at satisfying some superordiante goal).

To get a useful stabilization effect, the C->C loop gain through
the organism has to be significant w.r.t. the disturbances effecting
C. In principle, C->C loop gains could achieve oscillations or
chaos, but is there ever selection for this happening in a loop
going through the environment? At any rate, I'd want to say that
a `control system' was an organism internal setup that has
been selected (via evolution or reorganization) for because of
its ability to stabilize a condition C via negative loop gain.

Atlhough I->I loop gains presumably always exist, it seems to me that
C->C ones needn't. For example, C might be `be undetected by bats'.
C is *not perceived* by moths. Instead (perhaps), moths perceive
evidence of the presence of a hunting bat, and if they can't manage
to zero that perception, do something else, which makes them look
like a falling leaf. So C, which is important for survival, manages
to get maintained without being perceived.

Condition Maintenance via control and S-R chaining does not look very
exhausitive to me--I wonder what the other possibilities are?

Avery.Andrews@anu.edu.au

Date: Sat Mar 28, 1992 10:51 am PST
Subject: Another try

[From Rick Marken (920327 10:00)

Here is another try at sending this post. Hope it works.

Greg Williams (920325-3) says:

>Then let's start a new discussion about how organisms are organized
>internally -- specifically, about whether certain actions are due to "chains"
>or to "loops." By "chains," I mean that INTERNALLY there are no perceptual
>signals going to comparators, resulting in error signals which influence the
>actions' "trajectories" during those trajectories; generally, I would expect
>"chains" to look like precalibrated "responses" to trigger-like "stimuli".
>Such "chains" could be "output functions" of loops, and all I think would be
>sacrificed in PCT is the idea that control is always continuous.

I think the first step in figuring out how organisms are organized internally
(I mean functional organization -- anatomical organization can be determined
by inspection) is to determine what kind of behavior that organization is
presumed to accomplish. That is why I think that a discussion of how
organisms are organized should begin with testing for controlled variables.
That is, it should start with data. I think we have precious little data
regarding the variables that organisms control. But we do have a lot of
what I would argue is misleading data which appears to show relation-
ships between various input variables and various output variables. If
we know that there are likely to be strong negative feedback relationships
in organismic behavior, then I think it behooves us to test to see whether
any behavioral variable is a component of such a loop or not. This is what
the test for the controlled variable does. If a behavior (such as the now
mythical moth fall) is, indeed, found to be uncontrolled, then a model
such as the one you describe above may be appropriate; and be able to

generate this particular behavior in a realistic environment.

But given the great misconceptinos that can result from ignoring the POSSIBLE existence (I would say, ubiquitous existence) of strong negative feedback in the relationship between organisms and their environments, I would say that it would always be prudent to assume that organismic outputs are part of a control loop -- and probably controlled variables themselves. If "the test" reveals no evidence that a variable is controlled or part of a control loop, then you can take any observed relationship between this variable and other variables at face value -- and start modelling the underlying mechanism to account for this relationship. If, however, the variable is controlled or part of a control loop then efforts to model the relationship must be informed by an understanding of how the negative feedback loop affects the appearance of these relationships.

So the first step in any attempt to build models of the INTERNAL functions of an organism should begin with a very clear understanding of the phenomenon to be modelled. Most models of the functional INTERNALs of organisms are based on the assumption that the S-R or input output relationships that are observed are just what they seem -- input-output relationships. But because organisms are locked in a negative feedback relationship with their environments, these relation-ships may be (MAY BE -- but I think almost always ARE) the observed responses to disturbances of controlled variables. If this is the case, then an S-R type model of the INTERNALs would be wrong because it is trying to explain a causal chain that doesn't happen to exist in that case.

And so we come full circle (how approapriate) to "Marken's Law" -- PHENOMENA FIRST; models second. We need the data before we start arguing about how to account for these data; obviously. The contribution of PCT is to show that behavioral data must not be taken at face value -- because the phenomenon you are dealing with might be (MIGHT BE-- and, I have a strong hunch, almost always is) the phenomenon of CONTROL.

Regards Rick

Date: Sat Mar 28, 1992 10:51 am PST
Subject: Believe it or not

[Rick Marken (920328 8:30)]

Well, nothing seems to be getting from me out to the net, and nothing is getting in to me from the net. I'm going open loop here. But in the spirit of showing that open loop behavior can happen, I will generate a brief, open loop response.

Gary Cziko (920327) asks:

>Is this what "motor control" researchers actually believe we do--throw our
>limbs around with pre-computed initial forces and then see where they end
>up?

Yes. Strange but true.

>This seems so obviously absurd I can't believe anybody would hold such
>a view.

That's why I think there is an agenda deeper that "understanding human nature" at work here. The idea that negative feedback might be fundementally involved in all behavior is a real disturbance to these people.(They will allow that feedback might be necessary to do some minimal tidying up of behavior -- but the main thrust of all research on how organisms behave is to find evidence for and explanations in terms of open loop systems. If PCT is a religion then conventional psychologists are members of the rival church).

> If people put forth such purely
>open-loop models, wouldn't it be very easy to show how such a model can't
>work due to the laws of physics?

Well, yes. But if nobody is listening then all the demonstrating in the world won't make much difference. And these people are sometimes pretty clever about inventing "red herrings" that seem like reasonable responses to these demonstrations of the problems with their models (they rarely build working models anyway, by the way -- at least, not models that work in real environments). The "motor control" area does seems like the place where PCT can administer the most concrete coup d'grace to "open loop" models of purposive behavior. But these "open loop" theorists are controlling for a principle. Maybe it's "science abhors teleology" maybe its "physical models are REAL science"). but they are controlling -- and control systems don't just revise their references (in this case, for principles) when there is a disturbance; they PUSH BACK -- and these people are VERY effective at pushing back. The content of their pushes may seem absurd to you (and to me) -- but they work for them (and their audience -- which is controlling for the same principle). It's actually a very interesting phenomenon in itself (if it weren't also so disturbing to some of my own principles).

Regards   Rick

Date: Sat Mar 28, 1992 10:51 am PST
Subject: Language

RM79/[From Bill Powers (920327.1800)]

Bruce Nevin (920327.1220) --

The linguistics thread is still looking good to me.

It's probably been lost in the mists of time and tangles of verbiage, but some time ago I proposed that there are two modes of control going on in language at the same time: control for communication of meaning, and control for conformity to linguistic rules. These two modes are not necessarily dependent on each other, but as you've said from time to time they must have some strong interactions -- how strong depending to some extent, I would guess, on formal education. Haynes Johnson, on Washington Week in Review, often gets into a mode in which he gets his meanings across by a series of disjointed phrases or even single words that are not any kind of recognizeable sentence, although you end up knowing what he means. Charlie McDowell, on the other hand, speaks in polished sentences that could be written down and sent to the typesetter. Johnson is a writer of note, and certainly can put out polished language when he wants to, even when speaking. But it's clear that there's no one universal connection between one's style of communicating meaning and one's adherence to grammatical conventions.

Your extensions of my initial diagram of sentence expansion are quite acceptable to me. I wasn't intending to "rule out rules" by my diagram -- only to say that it didn't seem adequate to me to have language doing the work of expansion when we commonly do the expanding in terms of meanings, directly. After the meaning-level expansion I suggested, I meant to say, but didn't, that the expanded meanings are then described, and that in the process of description the adopted rules of sentence construction as well as those of indicating meanings come into play. As you have retained my suggested meaning level of expansion, there seems to be no problem here, unless what I just said has created one.

A thought that's been in the back of my mind for a long time seems a little closer now to being expressible. You've complained occasionally that I seem to be demanding the meanings themselves from you in cases where, as you say, you have to use words to communicate them. What's been bothering me can be illustrated by the verbal expansions. I've taken it for granted that these expansions are produced (in principle if not in the heat of generating a post) from a formal system based strictly on word-occurrances and empirical judgments of the likelihood of various usages. But the actual expansions I've seen don't seem to have this formal character. I would expect, from a formal system, that an input string of words would lead to _some particular expansion_ that anyone applying the same formal system would come up with: the VERY SAME expansion, no matter who applied the rule. But the offered expansions, and revisions and alternatives mentioned at the time and later, seem to indicate that the expansion rule is not totally formal -- in fact, that it can't be applied without referring to meanings.

In a recent post you said something that clarified this for me. In describing how you produce an expansion, you mentioned that one criterion was that the expansion HAD TO HAVE THE SAME MEANING as the reduced form or intermediate forms. A formal rule-based and empirical principle of expansion would not require that -- as I had been (mis) understanding what you claimed, it would produce sentences with the right meanings automatically.

Now I understand that the formal system can at best introduce constraints: if THAT's what you want the sentence to mean, then THIS is the way you have to expand it according the conventions of English.

It seems to me that the problem has now shifted to a subject that we haven't discussed, but which I've mentioned in passing now and then: how to analyze the process we call "description." This is the process of turning an experience into a phrase that means that experience. Or perhaps less directionally, it's the problem of deciding whether a given phrase is in fact a description of a given experience, and if it isn't, deciding what would make it a better description. A valid expansion not only has to obey the socially-agreed or learned rules of dependency and so on, but it has to be a valid description of an expanded meaning (for incomplete meanings, I would agree that the conventional expansions can also suggest missing meanings).

Behind all this I am still trying to unearth mental processes that are being used by linguists in the application of their principles, but which are taken for granted and not included in those principles. I would be asking similar questions of a mathematician in the course of finding out how Ax + Bx = C is transformed into x = C/(A+B). The mathematician would, like the linguist, begin by telling me the theorems that justify the transformation, which in this case I would already know. What I would be asking the mathematician would not be what the rules are that justify the transformation, but HOW THE MATHEMATICIAN GOES ABOUT APPLYING THOSE RULES. This requires something more than just demonstrating how the rules are applied: it requires stepping back and trying to notice what it is to be a mathematician. To do this, the mathematician has to abandon, for the moment, the point of view from which it's obvious that the rule justifies the transformation, and to try to see what "justification" is, what a "rule" is.

When you said that you make reference to meanings while constructing an expansion, you were telling me something about being a linguist. This, I think, is the only way we can find out what perceptual processes and what control processes underlie the linguistic methods that we try to describe in words. The words themselves are the product of these processes; they are not doing the work. But by stepping back and watching how the words are manipulated, it may be possible to see beneath the surface and get a hint as to the nature of the underlying processes.

>I'm not sure you realize that construal of "sentences expressing
>background knowledge" are not among the expansions I am talking about.

Yes, I realized it, but just ignored that subject. I'm still thinking of simpler situations in which the necessary background is in the foreground, as it were, and no recourse to special sublanguages is involved. I hope it's clear now that I'm not denying the linguistic process that parallels the nonlinguist one. As if any such denial from this quarter would carry much weight. You're really awfully good to treat my uninformed commentaries as if they were important.

By the way, the correct diagram for how a control system achieves food pellets is

```
                            /   variable   \
Stimulus--> food pellets ->   pressings
               ^^           \    of bar     /
               ||                   ||
               ||                   ||
                 ===================
```

The stimulus is something that disturbs the availability of food pellets. The bar-pressings also affect their availability.

Avery Andrews (920327) --

A negative goal is hard to handle conceptually and in a model -- "I don't want to see a unicorn." Philosophers seem to make a lot of the fact that in order to state a negative you implicitly have to recognize it as existent in some way. "There's no such thing as a 20-pound mouse." As a what? "A twenty pound mouse." Oh. One of those.

When you cast the same ideas in terms of perceptions, it becomes easier (real perceptions or imagined ones). Any perception can exist on a scale from none of it to some maximum amount:

```
     0 (No lion) --------------------------------------lion
```

Between the extremes you're perceiving more than no lion at all, but less than the maximum amount of lion-ness. When I'm in a zoo, my reference level for perceiving lionness can safely be set rather high:

```
     0 (No lion) --------------------------------------lion
                                                    ^
                                               ref amount
```

but if I'm in an African national park where lions roam free, I will set a considerably lower reference level for the same perception:

```
         No lion --------------------------------------lion
                    ^
                ref amount
```

Notice that I don't set the reference amount to zero. Lions are dangerous and they creep around through tall grass. It is much better to see just a little lion-ness (a tiny image far away) than none at all, because if you don't see any at all you don't know where the lion is. Same principle for poison ivy. You want to see the poison ivy, but not right up close, like underfoot. You're shown what it looks like so you can avoid it, which is paradoxical until you think in terms of high and low reference levels for a specific perception.

Avoidance means setting a low or zero reference level for a given perception. the perception itself is always defined positively. Speaking this way, you don't have to mention a perception and at the same time indicate that you don't want it to exist. You can say, I can imagine a unicorn but I have little desire to treat it as a factual being. My reference level for the proposition "Unicorns exist" is zero, or not very much. Speaking of factual beliefs as goals may sound odd, but think of "black people are inferior."

Try this on. You have a goal for children to live, and one means is to set your reference level for percieving them in danger (what you perceive as danger) to a very low level. Not necessarily zero -- you don't want them to grow up helpless -- but certainly not so high that they would get hurt.

Your conjectures seem to me to be an exploration of the logic or program level of perception and control. Even this level can deal in continuous variables.

Best to all,

Bill P.

Date: Sat Mar 28, 1992 10:59 am PST
Subject: Re: Open/Closed Loops

[From Rick Marken (920327 12:00)]

Greg Williams (920325-3) says:

>Then let's start a new discussion about how organisms are organized
>internally -- specifically, about whether certain actions are due to "chains"
>or to "loops." By "chains," I mean that INTERNALLY there are no perceptual
>signals going to comparators, resulting in error signals which influence the
>actions' "trajectories" during those trajectories; generally, I would expect
>"chains" to look like precalibrated "responses" to trigger-like "stimuli".
>Such "chains" could be "output functions" of loops, and all I think would be
>sacrificed in PCT is the idea that control is always continuous.

I think the first step in figuring out how organisms are organized internally (I mean functional organization -- anatomical organization can be determined by inspection) is to determine what kind of behavior that organization is presumed to accomplish. That is why I think that a discussion of how organisms are organized should begin with testing for controlled variables. That is, it should start with data. I think we have precious little data regarding the variables that organisms control. But we do have a lot of what I would argue is misleading data which appears to show relation- ships between various input variables and various output variables. If we know that there are likely to be strong negative feedback relationships in organismic behavior, then I think it behooves us to test to see whether any behavioral variable is a component of such a loop or not. This is what the test for the controlled variable does. If a behavior (such as the now mythical moth fall) is, indeed, found to be uncontrolled, then a model such as the one you describe above may be appropriate; and be able to generate this particular behavior in a realistic environment.

But given the great misconceptinos that can result from ignoring the POSSIBLE existence (I would say, ubiquitous existence) of strong negative feedback in the relationship between organisms and their environments, I would say that it would always be prudent to assume that organismic outputs are part of a control loop -- and probably controlled variables themselves. If "the test" reveals no evidence that a variable is controlled or part of a control loop, then you can take any observed relationship between this variable and other variables at face value -- and start modelling the underlying mechanism to account for this relationship. If, however, the variable is controlled or part of a control loop then efforts to model the relationship must be informed by an understanding of how the negative feedback loop affects the appearance of these relationships.

So the first step in any attempt to build models of the INTERNAL functions of an organism should begin with a very clear understanding of the phenomenon to be modelled. Most models of the functional INTERNALs of organisms are based on the assumption that the S-R or input output relationships that are observed are just what they seem -- input-output relationships. But because organisms are locked in a negative feedback relationship with their environments, these relation- ships may be (MAY BE -- but I think almost always ARE) the observed responses to disturbances of controlled variables. If this is the case, then an S-R type model of the INTERNALs would be wrong because it is trying to explain a causal chain that doesn't happen to exist in that case.

And so we come full circle (how approapriate) to "Marken's Law" -- PHENOMENA FIRST; models second. We need the data before we start arguing about how to account for these data; obviously. The contribution of PCT is to show that behavioral data must not be taken at face value -- because the phenomenon you are dealing with might be (MIGHT BE-- and, I have a strong hunch, almost always is) the phenomenon of CONTROL.

Regards   Rick

Date: Sat Mar 28, 1992 11:00 am PST
Subject: Re: Nonliving Closed Loops; "Motor Control"

[Martin Taylor 920328 14:00]
(Gary Cziko 920327 0900)

>Recent posts from Rick Marken, Bill Powers, and Martin Taylor (too much
>trouble to look up all the dates) have led me to the conclusion that closed
>loops do exist in non-living, non-human-made situations, but that they
>don't have BOTH negative feedback and high gain. Is this right?

No. Bill makes that claim, but I think it is wrong. See below.

>A pendulum has negative feedback but the energy that goes into restoring
>the pendulum does not exceed the energy that went into disturbing it (I
>suppose the restoring energy has to be less because of energy loss along
>the way).

Pendulums are not the kind of system we are talking about. We are talking about systems with strong energy input and output. In other words flow-through systems. All negative feedback systems are of this kind (I think). The gain can be as high as the relative energy rates of the flow through and the control. Typically those ratios are very high, perhaps millions, whether or not the system is living. The higher the ratio, the more stable can be the structure that is maintained by the negative feedback.

> But only living
>systems or engineered systems made to simulate living systems have BOTH
>negative feedback and high loop gain.
>
>Is this on the right track? Let me know.
>

No. I think the correct statement is that only living systems or their designed simulations have control, which is the use of negative feedback to bring a perceptual input to a determined reference state.

Martin

Date: Sat Mar 28, 1992 11:10 am PST

Subject: More open loop comments onn open loop behavior

[Rick Marken (920328 11:00)]

I want to clarify the points I made on friday in response to Greg Williams suggestion that we direct our discussion to hypotheses about the INTERNAL organization that produces the observed behavior of living systems. As usual, Bill Powers makes the point I wanted to make -- and far more clearly and concisely -- in his 1978 Psych Review article (Quantitative analysis of purposive systems) that is reprinted in the the Living Control Systems collection.

The crucial point is on p. 146 of the book. Two equations are found at the top of the page:

q.o = 1/g[q.i* - h(q.d)]   (1) and

q.o = f[h(q.d)]   (2)

Both equations describe the functional relationship between an environmental (disturbance) variable, q.d, and a system output variable, q.o. In both equations, the function h() is the physical law that maps the distal stimulus variable to the proximal stimulus variable (q.i). If the proximal variable is visual then h() can be thought of as a linear multipler.

Equation 1 is the relationship between distal stimulus and output response for a system where there are strong negative feedback effects (from q.o to q.i -- the latter being the proximal stimulus). The functional relationship between stimulus, q.d, and response, q.o, is the inverse of the feedback function, g(), that relates output to proximal input (q.i = g(q.o)). This means that the observed relationship between input and output has nothing to do with characterisatics of the organism (the INTERNAL organization that we are trying to understand, presumably). If you do an experiment with a negative feedback system where you manipulate a stimulus (q.d) and measure a response (q.o) and then plot q.o as a function of q.d then the shape of that plot depends on the shape of the feedback function (g()) and not on the properties of the organism!!! Conventional psychologists do this kind of experiment to understand the internal organization of organisms. Equation 2 shows that they would be learning about the internal organization of the organism IF they were dealing with an open loop system (what Bill called a Z system -- one with zero negative feedback effects of its outputs on its inputs). Equation 2 describes the functional relationship between distal stimulus (q.d) and output (q.o) for a Z system. This relationship depends on the function f() which is the "organism function". f() is a description of how the nervous system of the organism transforms inputs into outputs. It is the function we must know if we are going to develop a model of the internals of the organism -- because it is the nature of these internals that presumably determines the nature of f().

But f() does not even show up in the equations relating inputs to outputs in a negative feedback system; it gets "cancelled out" in a sense by the feedback effects.

So this is really the point I was trying to make on friday. We will find input-output relationships when we study organisms. But if the organism happens to be in a negative feedback SITUATION with respect to the relevant inputs and outputs, then observed relationships between input and output tell us nothing about the organism, only about it's environment. Models based on input-output studies of negative feedback systems are, thus, models of the environment, not of the organism (although the components are implemented as internals to the organism part of the model; if the organism is actually controlling a visual variable, for example, your organism model -- which is a model of 1/g() -- is a neural net model of the inverse of the laws of optics).

So, before building models of organisms, we should first find out if the variables involved (paricularly q.i and q.o) are part of a negative feedback SITUATION. This is done by tesing for controlled variables.

Regards   Rick

Date: Sat Mar 28, 1992 12:32 pm PST
Subject: S-R theory its own self

[From Bill Powers (920328.1030)]

Rick Marken (920328) --

Yes, phenomena first. It's hard to decide what is the phenomenon without some kind of model in mind, and it's even harder to recognize when you're unconsciously using a model to define the phenomenon instead of simply looking at what you can in fact observe as opposed to imagine.

I think we need to clear out a lot of underbrush in this S-R vs. control argument. We keep getting hung up on "what kind of system is this?" when the answer to that question doesn't really matter. The basic problems can be defined without getting inside the organism at all -- that comes later, when you try to think of the sort of internal organization that could account for what is seen externally. The most important problems have to do with the basic naivete of the S-R approach. Prepare for a lecture (Rick, you don't have to listen).

Greg Williams said "Let's forget about comparators that result in error signals that influence the actions" of the system. This means not thinking in terms of a particular implementation of the observed relationships, but we can't ignore the phenomena that comparators, error signals, and output functions are intended to explain.

Without a model, we assume that there's some sort of overall response to an external variable. If the objectively measured state of the stimulus variable is s, then the response r will be,

$$r = A(s - s^*)$$

The s* simply defines that state of the stimulus at which the observed response is zero. This will depend in part on the measurement scale. We can say that shivering rate v is a function of skin temperature t, so

$$v = A(t - t^*).$$

Clearly, shivering doesn't begin either at 0 Celcius or 0 Fahrenheit, so we have to use t*, on the appropriate scale, to indicate the shivering threshold temperature in the units we're using. The rate of shivering grows as the temperature departs from the effective zero of temperature t*, and A describes how rapidly shivering increases as temperature increases. The observations show that A is a negative number, so that shivering will actually increase only as temperature decreases below t*. Also, observation shows that this equation only applies for $t < t^*$.

In either S-R theory or control theory, we could refer to s* as the reference level of the stimulus s. It is that point on the scale of measurement of the stimulus at which the response just becomes zero. That is the formal definition of a reference LEVEL in control theory, too: it is an observational definition. The control model merely proposes a mechanism for inserting a variable value for s* -- a "comparator" and a reference SIGNAL.

Note that for an on-off stimulus, measured as 0 or 1, there is still an s* -- it is still "that value on the scale of measurement of s at which the response becomes zero." If the response is to the presence of the stimulus, then s* is zero. If it is to the disappearance of the stimulus, then s* is 1.

There are types of stimuli and types of responses in S-R theory. With respect to stimuli, the least attention has been paid to continuous stimuli. A very common definition of a stimulus is as an event, an impulse taking place at one instant of time. Another one is "onset" or "offset", which is defined as the first derivative of a continuous stimulus variable undergoing a sudden change. Another employs a logical condition of some sort, for example stimulus greater than (or not greater than) some threshold value. Even more general concepts of a stimulus are used, such as the presence of some abstract condition near the organism, like a "threat." B. F. Skinner introduced rates of occurrance as a type of stimulus.

Responses likewise come in different types. The continuous response is again neglected in general. Skinner's rates of occurrance are used. Event responses are also common (first derivatives or changes). There is a latching response, where the occurrance of a stimulus turns the response on, and it then stays on regardless of further stimulation. A related type of response is the "triggered" response, where a stimulus initiates some series of actions that plays itself out regardless of further stimulation. There are also switching responses, where the stimulus occasions a change from one mode of action to another.

All of these modes of stimulation and response can be represented in the general S-R equation, $r = f(s-s^*)$.

A basic question about these kinds of S and R concerns their objective existence. Event-type stimuli and responses, for example, are almost always artifacts of definition or experimental conditions. A rat may move toward the entry of a maze and then start down it, all in one continuous flow, but the only recorded event is the instant its nose breaks a light-beam. On the way to the bait, the rat may sniff down side-alleys, pause, move slowly or rapidly, until it finally enters the goal box, sniffs the bait, and moves up to it and takes a bite. Only the instant of the bite is recorded as an event. When I was being taught experimental psychology, it was emphasized that experimental designs MUST provide for such measurements of critical events, or else there would be no events to measure! I am therefore suspicious about S-R relationships expressed in terms of events. I am equally suspicious of most other ways of defining stimuli and responses, for reasons that can be deduced from the rest of this lecture.

I have expanded the types of stimuli that can be considered, at least for human subjects: 11 categories of them.

In modeling behavior from any viewpoint, it's necessary to be much more careful about measures of either stimuli or responses than when making armchair judgments. The human observer all too easily and inappropriately projects human perceptions into the scene, especially for stimuli but almost equally for responses.

Take the cockroach's "escape" response to a "threatening movement." The words in quotes represent human interpretations of the situation derived with human senses and interpreted by a human brain. The human being can go on, and claim that the reason for this response is to promote survival of the cockroach. But we'll keep it simple.

A modeler can't take such human interpretations into consideration. To make a model, one has to provide the model itself with all the capacities needed to sense the external state of affairs. How do we equip a cockroach to detect a "threat?" This is impossible to do. Instead, we have to ask "What would be stimulating the cockroach under the conditions where the human observer sees a threat to the cockroach?" One answer is that the cockroach might detect moving air displaced by the approach of a large object. It might detect infrared radiation. It might detect light patterns on its retina. It might detect vibrations in the floor.

None of those stimuli, of course, is a "threat." Only the human being would associate them with such an abstract concept. The cockroach doesn't classify these stimuli as a human being does. Human logic isn't involved in the cockroach's response -- it doesn't think "every time stimuli of this kind have occurred, something bad has (nearly) happened." All such considerations are irrelevant to the cockroach. The cockroach does not respond because of a threat. It responds because of wind, infrared, light, or vibration.

Neither does the cockroach respond by escaping. There can be no "escape response" in a cockroach -- that is a human classification and has no relevance to the cockroach. The cockroach can move forward, backward, and sideways at various speeds, and that is all. None of those movements inherently constitutes "escape." The same movements are used in approaching food, avoiding obstacles, and seeking dark damp places.

Escape can't be defined strictly in terms of one of the actions available to the cockroach. It's a consequence, a relationship between the cockroach and something else, that's affected by the way the locomotive machinery is used. Whether these motions amount to "escape" depends partly on the movements and partly on what is going on independently in the environment. Escape, as seen by the human observer, is a very particular relationship between the movements of the cockroach and the spatial relationships between the cockroach and the "threat." The cockroach must move AWAY from the threat.

Now "away" is a peculiar word. It seems to define a direction, but in fact it refers to a relationship; the actual direction remains unspecified. In order for the cockroach to move reliably "away" from a "threat," it must somehow know WHERE THE THREAT IS. If we now look back at the proposed stimuli (wind, infrared,light, and vibration), we can see that they are inadequately defined. It is not wind, for example, that is the stimulus, but WIND FROM A PARTICULAR DIRECTION RELATIVE TO THE COCKROACH'S BODY. There is no sensory receptor that can report wind velocity AND DIRECTION. Directional information must somehow be derived from WHICH receptor is stimulated, not from the information carried in a particular sensory nerve fiber.

The escape response of locomoting in a specific direction, therefore, must be based on sensory information that indicates both the existence and the direction of the threat. The direction of the threat must translate into a coordinated set of leg movements that carry the cockroach away from the sensed direction of the threat. The velocity of movement, presumably, would depend on the intensity of the stimuli.

By looking carefully at the details, we can see how the cockroach might respond to various stimuli in a way that takes it in a direction away from the threat. Stimulation of specific points on its body would generate specific combinations of locomotive movements aimed toward the side opposite to the points of stimulation. Thus the cockroach would accomplish something a human being might classify as "escaping" (or at least "fleeing") from a "threat", without any stimulus that specifically means "threat" or any response that specifically means "escaping."

These details tell us that the characterization of actions as an "escape response" is naive, being loaded with subjective interpretations and hidden assumptions. By looking at the details in this way, we move much closer to appreciating the world of stimulus and response that is relevant to the cockroach, and away from the anthropocentric interpretation of that world. Even though no model of the cockroach has been proposed, we have approached the external situation in such a way that we could now begin modeling without having to figure out how to model such categorical concepts as "threats" or "escaping" inside a mere cockroach.

There is one final consideration. Having descended to the level of detailed specific stimulations, we are now in a much smaller world, and we ought to look at it on a much smaller time scale. To a human being, a puff of air is simply an event, treated as it if occurred at one instant. But now that we've shrunk down to the scale of the wind-detecting hairs on a cockroach's body, we can experience the puff of air as it really occurs. There is first a slight stir as molecules begin drifting coherently in one direction. The flow picks up, and the hair bends slightly. A neuron here and there passes its threshold of stimulation and little ticks begin, like popcorn just starting to pop. The flow of air becomes stronger and stronger; the hair leans more and more; the ticking increases to a rattle, then a buzz, then a roar. Then the flow begins to slacken, the popcorn pops more sporadically, and the wind gradually tails off to lower and lower levels until it is still again.

This is the event called a "puff of air." During this event, a number of responses were getting under way. At a certain level of buzzing of the sensory receptors, the locomotive machinery began to turn the body of the cockroach. This, of course, immediately began to alter the direction of wind-flow around the body, affecting the wind velocity at the sensor beside our vantage point. This modified the sensor signal, and the response, which had been aimed at one final state, is now aimed at a slightly different final state. All through this puff of air, the body is moving and the air flow is changing as a result.

We can now see that our general equation for a stimulus-response relationship, r = f(s-s*), is wrong. We have left out the effect of the response on the stimulus. The equation should read,

$$r = f(g(r) + s - s^*)$$

where g is now the function expressing the addition of response effects to the effective stimulus. While there's no general principle that says g(r) MUST be a nonzero function, including it in the equation allows the equation to cover those cases in which s is in fact affected by r.

In those cases where g(r) is not identically zero, the response affects the stimulus while the response is in progress. This is the situation defined as "feedback." There's no theory here. Either feedback is present or it's not. If it's present, the only remaining questions are "how much?" and "what's the sign?" The question is not "Should we take it into account?" There is no choice about that if we want a correct characterization of the S-R relationship.

Now we come to the crucial question in considering S-R theory. The question is, how common is it for g(r) to be other than zero? This question is not a theoretical one: it is a factual one. But we have to make sure to ask this question of the right data.

If we view an S-R relationship as an escape response to a threat, using naive human categorizations of the observable events and relationships, we will see no g(r). That is, if I stamp my foot near the cockroach, creating a threat, the cockroach's escape response will have no effect at all on the stamping of my foot. The conclusion would be that there is no feedback and this is a simple S-R phenomenon. That conclusion, as we have seen, would simply be wrong.

Even if if we used a microanemometer to measure the puff of air itself instead of the stamping of the foot, we would erroneously conclude that g(r) = 0. We have to measure the puff of air where the measurement is relevant: we must mount the anemometer on the cockroach's body next to the hair in question. The only relevant movement of air is the movement relative to the hair, in a frame of reference that is attached to the cockroach. When we make that measurement, we find that g(r) is no longer zero. There is feedback. In fact, the "escape response" drastically modifies this relative air velocity while the response is in progress.

Now we are looking at the TRUE stimulus, not just at a human observer's careless and anthropocentric evaluation of the surrounding environment. We are seeing the physical process to which the cockroach is actually responding, and it is quite different from the human observer's interpretation. We are now looking at the world that the cockroach experiences.

In all of this long lecture the only "model" I have proposed was the idea that sensory nerves response to stimuli. The rest has been concerned only with information available to an observer outside the behaving system. I am really talking about thoroughness and care in observing the details of behavior. I am talking about becoming aware of the way human observers carelessly impose their interpretations on global phenomena and fail to think in terms of details that are all-important to arriving at the right interpretation.

Control theory arises simply from looking carefully at the details of any stimulus-response situation. It is, if you like, stimulus-response theory done right.

I don't care in the least whether some responses are indeed unitary and some stimuli are instantaneous, whether g(r) is zero or nonzero, or what is found really to be the case. What I'm concerned with is getting away from the sloppy habits of observation that have led to S-R theory as it now exists, the projection of inappropriate kinds of interpretation onto the very act of taking data, so that the wrong processes are noted and the absolutely critical ones are glossed over as "mere detail." It's simply not possible to understand behavior correctly if you stand back and generalize about made-up variables having no proven relationship to the organism. The reason that behavioral science has come up with such terrible unpredictive uncertain results is not that behavior is that way, but that behavioral science is that way. And behavioral science is that way because of sheer sloppiness of observation.

I claim that in fact, g(r) is nonzero in essentially any kind of behavioral situation that can be found. Every response alters the very stimuli that lead to it immediately and strongly. The real stimuli, that is, not the ones seen through the abstractions of casual and subjective observation. Given my claim, the remainder of HPCT follows. Make any model you like of the organism's interior. But it must be able to operate when g(r) is other than zero.

Best,

Bill P.

Date: Sat Mar 28, 1992 4:31 pm PST
Subject: I listened anyway

[From Rick Marken (920328 16:10)]

Bill Powers (920328) ---

I listened to the lecture anyway. I'm finding that this discussion of SR is really helping to clarify my understanding of the relationship between control theory and other theories of behavior. I now understand better than I did before (I'm not sure that I did really understand it fully) that the FACT of negative feedback

just flat out eliminates your ability to SEE the actual functional
relationship between input and output. I mean I kind of knew it
was true -- but I didn't really understand HOW true. That is a
REALLY HEAVY FACT. The importance of your "cubic function" demo
in the Psych Review article is now clearer to me than it ever
was. Talk about psychology not knowing what hit it. They'll
go nuts if they ever figure that one out.

Regards Rick

Date: Sat Mar 28, 1992 10:38 pm PST
Subject: logic, negation

I don't see that much use for graded perception of lion-ness. Like, a little bit of lion seen through tall grass is much more unnerving than lots of lion seen through nice thick iron bars. I think that what the continuous variable ought to be is estimated likelihood of bad things happening (ranging for 0 to 1), and conjecture that something sort of like the logical deductions I exhibited might be involved in calculating them, along the lines, perhaps of an assignment of values chosen from {1,0} for intentions (1 = do it, 0 = don't do it) that minimizes a bad consequence score.

One consideration that urges me to think that way is the fact that people are obviously quite bad at predicting bad consequences of novel kinds of courses of action, even when the logical deduction from the action to the consequences is very short, and seems shockingly obvious in retrospect. This is, I believe, because we actually maintain long lists of bad things that we don't want to happen, and try to figure out how likely they are to, given how things are. Children don't have such large lists, which is one reason they do so many stupid things.

As for philosophers and negation, I think I can deal with that sort of trouble (the other day I presented a quick story about how to do `Zeus doesn't exist' and didn't get shot down--I haven't been hanging around with these guys for three months for nothing).

Avery.Andrews@anu.edu.au

Date: Sun Mar 29, 1992 4:45 am PST
Subject: The details

From Greg Williams (920329)

>Bill Powers (920328.1030)

A very nice post, with which I agree almost completely. By that, I mean that I agree with the basic, underlying, important points. As you note, the DETAILS are all important. I have one question about the details of your cockroach "escape" example.

>Even if if we used a microanemometer to measure the puff of air itself
>instead of the stamping of the foot, we would erroneously conclude that

>g(r) = 0. We have to measure the puff of air where the measurement is
>relevant: we must mount the anemometer on the cockroach's body next to the
>hair in question. The only relevant movement of air is the movement
>relative to the hair, in a frame of reference that is attached to the
>cockroach. When we make that measurement, we find that g(r) is no longer
>zero. There is feedback. In fact, the "escape response" drastically
>modifies this relative air velocity while the response is in progress.

Are you claiming that there is actual data showing that (a) in the case of short-duration air puffs, the body orientation changes before the air puff is over, and/or that (b) if the body orientation does change before the (say, somewhat longer-duration) air puff is over, the movement of the cockroach is actually influenced by the change in alignment between the body and the continuing air puff? Or are you hypothesizing that feedback would be shown in this example if someone did the appropriate experiments (which might or might not actually have been done)? Or are you just making a rhetorical point and postulating that feedback is there in your exemplar case to help make the point?

>I don't care in the least whether some responses are indeed unitary and
>some stimuli are instantaneous, whether g(r) is zero or nonzero, or what is
>found really to be the case. What I'm concerned with is getting away from
>the sloppy habits of observation that have led to S-R theory as it now
>exists, the projection of inappropriate kinds of interpretation onto the
>very act of taking data, so that the wrong processes are noted and the
>absolutely critical ones are glossed over as "mere detail."

Yes, yes, yes! Everyone needs to "ask the organisms," not themselves!!

>I claim that in fact, g(r) is nonzero in essentially any kind of behavioral
>situation that can be found. Every response alters the very stimuli that
>lead to it immediately and strongly. The real stimuli, that is, not the
>ones seen through the abstractions of casual and subjective observation.
>Given my claim, the remainder of HPCT follows. Make any model you like of
>the organism's interior. But it must be able to operate when g(r) is other
>than zero.

There ARE a lot of cases where g(r) is demonstrably nonzero. But it is quite a leap to your general claim. I wish you could see that the worth of HCPT modelling does not hang on the truth of that generality. In fact, I think leaving the question of the truth of your generalization open, at least for now, would boost the credibility of HPCT ideas for many behavioral scientists. It is a matter of admitting humility until the data are in. That humility in no way compromises the significance of HPCT in explaining organismic CONTROL, when it is actually found to be present.

Greg

Date: Sun Mar 29, 1992 10:54 am PST
Subject: Re: Lots of peevish remarks

[Martin Taylor 920329 13:30]
(Bill Powers 920327 07:30)
>
>
>I have a lot of trouble relating "maintaining entropy" to anything real,

>like standing up or pointing to a target (or even trying to write good
>sentences). I don't think that I need to maintain entropy in order to do
>anything, or that if I do something, I do it by means of maintaining
>entropy. Generalizations like this don't substitute for an explanatory
>model.
>
>Do you have any numbers for what "quite strong gains" means? Are you saying
>that the gains are high enough to prevent disturbances from significantly
>distorting a "far-from-equilibrium energy flow?"

I'm not making an explanatory model of anything when I refer to the maintenance of entropy. I'm describing negative feedback. Environmental disturbance increases the entropy of the system, and negative feedback reduces it. In a stable controlled system the entropy remains stable despite disturbance. It's a statement of fact that gives you a different slant on what happens in a negative feedback system. In a positive feedback system the reverse happens. In a high energy flow, the system can export entropy ("garbage" if you like) to maintain its stability. That's where the stable self-organized structures come from.

Yes, I mean that disturbances, within limits, will be resisted in such a way that the disturbed structure returns to its undisturbed state. Naturally, there will be disturbances too strong for restoration, but that happens with living systems as well.

I maintain that the distinction you are looking for is not the high-gain negative feedback, but the reference-signal + comparator structure.
==============
I agree about high-level errors being detected rapidly on occasion. In the Layered Protocol theory of communication, we call this "early interpretation," and take it to be an essential aspect of effective communication, just as it is in normal (i.e. dealing with the non-living world) behaviour. This doesn't alter any of my comments about the need for statistical interpretation, which I think you misunderstood. I am talking about statistics within the perceptual input structure, NOT statistics of experimental results. That's a different issue, about which we have a minor disagreement.

There's no need to get into a "yes it is" "No it isn't" argument. That's fruitless. What I hope to do, probably starting when I get back from my trip in June, is to come to a resolution of the problems of the thermodynamics of the control system hierarchy. It is for that reason that I commended the granularity article in science.
============

>I think that the normal (non-pathological) state of the hierarchy at every
>level is one in which perception very nearly matches reference at all
>times. Only the little variations in error cause shifts of reference
>signals at lower levels, and the perceptions at those lower levels simply
>track the perceptions, with a subjective lag that's zero. This is why
>control -- that is, strong negative feedback -- has gone mostly
>undiscovered. We refer to these processes as "doing."

OK, except for the caveat that a subjective lag of zero is not the same as a loop delay of zero. The latter refers to how long it takes the perceiving system to acquire information about the effect of action in the presence of all sorts of environmental disturbances and disturbances induced by competing ECSs.

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

>If we take these observable relationships into account as we look at models
>for internal chains (there are no _internal_ "loops"), I will be quite
>content. Only you do it this time: I've already done it.
>
Why do you say there are no "internal" loops? Do you mean that there are
no loops within a single ECS? Do you mean that imagination doesn't happen?
Or that kinaesthesis doesn't happen? Or are you making the assertion, that
I expect to take issue with later, that there are no cross-connections
within a level of the hierarchy?
>--
>Jeff Dooley (again)
>
>Reorganization doesn't require "chaotic" behavior in the technical sense.
>All it needs is random change. Freeman tells a wonderful story, but it's
>practically all rhetoric. He could be right. Or he could just be making up
>stories. I can't tell.

No, it isn't reorganization that requires chaotic behaviour. It's not at all clear that reorganization even has the basic characteristics that could lead to chaos. The chaotic behaviour (or, more probably, near critical behaviour) is required for the perceptual functions, and it is to make that behaviour useful that we have categories. I think that's what Freeman is saying, as well. But even if he isn't, I am.

Martin

Date: Sun Mar 29, 1992 2:17 pm PST
Subject: Category, configuration, sequence control

[From Bill Powers (920329.1400)]

Avery Andrews (920329) --

I think we're talking about differences in style here:

>I don't see that much use for graded perception of lion-ness. Like, a
>little bit of lion seen through tall grass is much more unnerving than
>lots of lion seen through nice thick iron bars.

If a person is used to dealing in either-or caterories and treating them as logical variables, then the reaction would be "Yipes, a lion, I'm getting out of here!" whether the lion was seen in a nearby thicket, behind bars, or snoozing 200 yards away. But you sort of slipped sideways from my point, because on a graded scale, you can pick HOW MUCH of the perception you want to experience. If you set your reference perception to zero regardless of the circumstances or the nearness of the lion, then of course you'll react maximally in all cases. But I'm saying that in the zoo, you can decide that it's OK to set a relative high/large/near reference signal for the perception of the lion, while out in the wild, you're likely to want to keep that perception weak/small/far. Of course if the only two cases you can perceive are lion or no lion, then you can't do this. But I think everyone really can perceive lion proximity on a graded scale, and suit the reference level to the situation.

What you say about estimating probabilities may apply in such a case. Actually, I don't think that people do much estimating of probabilities or calculating of payoff matrices, although cognitive theorists certainly do. I don't think people do much predicting, either -- it's much easier to handle "predicting" in terms of reference signals and imagination. Remember that all during the time these digital-like concepts of the brain as a rational logical computer were developing, everyone thought that real-time purpose and actual goal-direction were figments of the mystical imagination. One common substitute for actual goals has been "outcome prediction." With control theory, that substitute isn't necessary -- we can just accept the reality of purposes and goals. Even a line of thinking that is well-developed and widely accepted isn't necessarily leading anywhere. I feel that a lot of concepts currently in use are just part of the whole "computer revolution" that got everyone off on the wrong foot in thinking about the brain.

The poison ivy example is probably better than the lion example where you have only partial control of nearness to the lion. If you're walking through the woods, you want to avoid contact with poison ivy, but at the same time you don't want to miss seeing it if it's there. Your reference level for seeing poison ivy is non-zero, but you don't want to see it up close. I think that in most cases like this people control for perceptions on a continuum, not categorically.

In general categorical control, literally carried out, is pretty poor. I know that people use it, but it really doesn't do them much good, or as much good as controlling the same variable in a continuum.

Dag Forssell has a nice example taken, I think, from the Deming approach to "Total quality management." In America, quality control is often done categorically: go or no-go. An error circle is set up, and the goal is to keep the measurements "within specs" -- inside the circle. Under the Deming approach, a target is set at the center of the circle, and the object is to get the measurements as close to the target as possible. There is a great improvement of quality in the latter case.

Greg Williams --

We may as well go public with the arm-model discussions. I received the paper by Atkeson [sic) and Hollerbach, "Kinematic features of unrestrained vertical arm movements" (Journal opf Neuroscience, Vol. 5, p. 2318-2330, Sept. 1985). Thanks for your usual helpfulness.

I see what you mean by the outward curvature on upstrokes and downstrokes, although there is a lot more variation between subjects than was my impression from what you said. I wish we had the original data -- it's hard to get any quantitative measurements of what the two joints were doing from the figures. There's also a critical piece of missing data: in these plots of visual-motor behavior, the position of the eye isn't shown, and it's not mentioned whether the head bent forward as the arm descended! The curvatures are not along shoulder-centered circles, or as near as I can estimate, eye-centered circles, although they might approach being eye-centered if the head nodded up and down during movements.

The authors also mention that the hand didn't maintain exactly the same relation to the wrist during the movements. Since the fingertip, not the wrist, was brought to the target, this puts some uncertainty into the data.

How much is hard to estimate. My model, of course, has only two joints, not three, in the vertical plane.

The authors don't mention whether the shoulder joint was fixed in space. The LED closest to the shoulder does seem to move, but it's hard to estimate where the center of curvature is, or whether it remains fixed. This, of course, would add two more degrees of freedom to the arm control, which I can't reproduce in my model.

The most interesting problem is the speed of movement and the shape of the tangential velocity curve. When the traces of tangential velocity are normalized for duration and amplitude, they all have very nearly the same shape. This looks like gain control. In my model, there's no provision for controlling speed of movements. Perhaps, by putting gain control into the visual system, this effect can be reproduced. This would only be germane for the upper range of speeds, however. If you're asked to draw a straight line from the starting position to the ending position, you probably can do it pretty well if you can go slowly enough. In fact you could draw a wiggly line, a semicircle, two straight lines with a bend in the middle, and so on. It's very difficult to separate higher-level control from the basic arm-positioning and target-tracking systems. In the authors' experiments, they gave no instructions as to what path was to be followed (of course they still assumed that the path was "planned"). As a result, we don't know whether the observed path was one the subjects intended to follow. Is that what we're seeing? If you introduce variations in the reference signals to the visual systems in my model, you can create any path you like.

The curvature problem is not so interesting. If the traces of fingertip movement were drawn with a line one centimeter wide, the difference from my model's behavior would look a lot smaller, particularly if you merged the data for all the subjects. For my part, if the position of the fingertip stays within a centimeter of the average real fingertip throughout a movement I'd be satisfied. This model has only two levels in it, and no correction for distortions at all.

By the way, whatever errors there are, they're not in the kinesthetic levels. Those levels will make the joint angles follow the reference signals in only a tenth of a second or so. The curvature problems aren't arising at that level, as you can tell by going to the imagination mode (dynamics off). The detailed path is determined by the visual systems, not by the kinesthetic ones and not by arm kinematics.

Martin Taylor --

Thinking generally about your problem of sequence control, I've had a thought that may be useful. It's similar to a thought I've had about configurations, so I'll start with that.

While we refer to a "configuration level" and call something like a human face a configuration perception, what goes on at that level must be more detailed than simply perceiving "a configuration." A face can change in a lot of ways and still be recognized as a face and not a hand (if not the SAME face). At the configuration level, there is probably a collection of attributes that make up configuration-space. These would be attributes like size, orientation, relative position, elongation and squashing, and so on. Perceptions in all of these dimensions together add up to what we call, for convenience, "configuration." But the configuration level must really be a multi-dimensional space (like the sensation level) in which particular configurations are represented as vectors with particular directions in this space. A configuration signal then indicates by its magnitude the magnitude of configuration perceptions projected onto the vector associated with a particular input function. Maybe this will help with your nagging sense that multiple dimensions have to get into the picture somehow.

Not quite as clearly, we can try the same idea on for sequences. "Sequence" is a name for a perceptual space. The attributes of sequences make up the independent dimensions of this space, and a particular input function defines a vector in this space (or some higher-level mathematical construct). The magnitude of the perceptual signal indicates the magnitude of the perceived sequence as projected onto this vector.

The attributes of sequences would include such things as ascending-descending (in any measure), closed or open, and whatever else you can think of. The perceptual signals wouldn't just indicate sequenceness; they'd indicate a particular combination of the attributes we perceive in sequences or orderings. Maybe you can think up some more of these attributes. "Alphabetic" ordering would be one: "ABCDEFHGIJK...". This example shows ALMOST aphabetic ordering.

This basic idea is probably going to help in defining other levels, too -- that is, considering the label for the level as indicating a perceptual space, with particular perceptual functions defining vectors in that space, and the dimensions being identified as possible attributes within that space.

This is making a lot of sense right now -- how is it coming through at the other end?

Best to all Bill P.

Date: Sun Mar 29, 1992 6:51 pm PST
Subject: PCT TV Programs

from Ed Ford (920329.19:47)

Gary - concerning your comments about my PBS program. I was working under two constraints, time and simplicity. Both segments were to equal 45 minutes to give the station time for their promotional stuff. Ultimately it was 46 minutes. I had control over the credits, and they ran exactly what I wanted but they could have added a second or two more. The commercial copy should have my publishing company info. During the two weeks of practice sessions, I kept trying to simplify the content while trying to maintain the integrity of control theory. The purpose for the credits at the end was first to give credibility to to the show's content (1st half) and control theory and second, references for those interested. The response to the program met the station's goals, and they plan to promote it to other PBS stations in June at some kind of gathering. If it's popularity continues to rise, then I plan to do at least two more, one on Freedom From Stress and one on Teaching Responsibility. The format MUST involve an audience and its participation, but there will be role plays, etc.

This brings up the intent of all my programs. It seems that when a person passes the "aha" threshold of understanding control theory, and I mean really understanding it, one's perception of others changes, quite radically. You begin to realize, slowly, that everyone you meet is truly a living control system, with all kinds of complex and inter-related agenda's and created perceptions, all of which you'll really never understand, cause you haven't had the exact experiences from which they've created their own unique perceptions and the creation of a perception from the same experiences could be quite different.

The difficult art of understanding and then learning to live with one of these systems as well as the potential satisfaction that comes from successfully living with another, that's what my first tape was all about. Understanding PCT helps you cut through all the baloney and miscellaneous crap that is so prevalent in today's self-help culture. It forces the spouse/parent/teacher/therapist to respect the unknown world of those with whom they deal. It teaches you not to push, not to judge, not to control or you'll do violence to their systems. Rather, you learn to respect, to ask, to teach the other person how to establish and evaluate their goals, priorities, standards, what they want, how they presently perceive things. PCT forms the basis from which you can help others build efficiency into how they operate their system, thus helping others to function more effectively so they can satisfy their own internal goals.

I really believe that when you really understand PCT, it allows for only one effective alternative when dealing with another; you've got to respect their system. As I read my CSGnet mail, I'm overwhelmed at times by the complicated stuff that's discussed. I've also notice the respect everyone has for others, which is the true mark of a control theorist. Once you understand the theory, you realize this respect is the only way to learn from and deal with another. I remember our first CSG meeting in Chicago, and the respect I was shown, even though (as I now look back) I knew precious little about control theory. It was that respect that told me that I was in the right place.

My present effort is to look for ways to help others who are struggling with their lives by getting them to look at what they are doing, not to try to control, push, or analyze them, but to teach them how to create a more satisfying life through a solid understanding of how our system operates, based on PCT. Basic to all this is the continuing respect I must maintain for the living control system with whom I am working. I realize this is all elementary to most of you, but it is what I'm trying to do.

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

Date: Sun Mar 29, 1992 6:58 pm PST
Subject: mpd

To: Mark Olson and others interested in MPD
From: David Goldstein
Subject: MPD & HPCT
Date: 03/29/92

On March 27, 1992, Mark Olson wrote about a friend with MPD. He threw out several questions to the clinicians in the group concerning MPD and HPCT. I have one patient with MPD, have read a few books on the topic, have attended a few workshops, and monthly participate in a study group made up of about six clinicinas who have patients with MPD.

Here is a book I would recommend to learn more about the subject:
Putnam, F. W. (1989). Diagnosis and treatment of multiple personality disorder. New York: Guilford.

Question: Is it OK to equate self-concept with personality? This is the way that Dick Robertson and I think of it within HPCT. The self-image is a perception at the systems level which is a controlled variable. A self-image perception is formed from combining lower level perceptions, with an emphasis on the principle level perceptions.

Question: What keeps this from being more prevalent? Is it only random variation and selection probabilities which make most of us have one personality? Something like 95% of people with MPD have a background of severe physical, sexual or psychological abuse. They also have a special talent to experience hypnotic phenomena. MPD is thought to be on the extreme end of a continuum of dissociative disorders which you can read about in the DSM-3-R.

A second point is that I am not so certain that we all have one personality defined as self-image. In some research I am doing now, it seems that "normal" people have more than one self-image. I am using Q methodology to carry out these clinical case studies.

Question: Is MPD "simply" a multiplicity of self-concepts residing at the systems level? That would be my guess at this point. Obviously, there are many differences between a normal person with multiple self-images and a person with MPD. But I don't think it is simple one versus many.

Question: If person B wants A not to see something (say, three words in a amessage) she can make A not "see" those words. How might this fit into PCT? A "normal" person with high hypnotic ability can demonstrate the above phenomenon which is called negative hallucination. I am not really sure how to model it. Is the perceptual signal wiped out? Is a person's awareness diverted from tuning into the perceptual signal? I don't know.

Mark, one word of caution with your friend. It is not unusual for one of the "alters", as they are sometimes called, to be suicidal or homocidal. It probably is wise to ask permission of your friend to meet with the therapist who can provide some concrete suggestions for you.

Best regard,
David Goldstein
internet: goldstein@saturn.glassboro.edu

Date: Sun Mar 29, 1992 6:58 pm PST
Subject: Modeling/Testing

[From Rick Marken (920329 18:40)]
a
b
c
Doesn't anyone want to comment on my "Behavioral Illusion"
post? Doesn't that strike some of you as a remarkable
discovery of Bill's? If there is strong negative feedback then
stimulus response laws are environmental laws-- not organism laws.
I would be especially interested in hearing from behavioral
scientists.

Greg Williams (920329) says to Bill Powers:

>There ARE a lot of cases where g(r) is demonstrably nonzero. But it is quite a
>leap to your general claim. I wish you could see that the worth of HCPT
>modelling does not hang on the truth of that generality. In fact, I think
>leaving the question of the truth of your generalization open, at least for
>now, would boost the credibility of HPCT ideas for many behavioral scientists.
>It is a matter of admitting humility until the data are in. That humility in
>no way compromises the significance of HPCT in explaining organismic CONTROL,
>when it is actually found to be present.

I agree that HPCT is valuable even if the general claim
(that g(r) is always non-zero) is not true. But I don't think
that claim is what affects the credibility of HPCT. It is, rather,
the implicit assumption on the part of most researchers that g(r)
is generally zero -- and that HPCT is unnecessary. What these
researchers refuse to accept is the possibility that g(r)
MIGHT be large and negative. Because of this, they see
no need to test for controlled variables. They simply
take their observations at face value and proceed to develop
models that are based on the assumption that they are observing
the behavior of a g(r) = zero system (a Z system). All PCTers
are saying is that g(r) MIGHT be non-zero and, if it is, then
they are up shit's creek without a paddle. But they ignore this.
I don't think PCT people are making grandiose claims; we
are just being ignored. I am perfectly willing to accept the
fact that g(r) is always zero. I would be surprised by such a result but
I don't want to prejudge. I say "just test for controlled variables";
that's all I ask of the people who are developing the "open loop"
models. If it turns out that the output variables that they are
dealing with are, indeed, generated "open loop" then their models
are just fine and I'll be happy; at least they eliminated
the possibility that they are dealing with a negative feedback
control loop by doing the test.

Perhaps this message has not been made clearly to the "other side".
The message is this: "We (PCTers) have no theoretical ax to grind.
We have a theory that explains control and we have used it successfully
to model behavior where control is demonstrably involved. We know
that if you are dealing with behavior that is part of a negative

feedback loop then observed input-output relationships are
mirrors of environmental (not organismic) functions. So we
heartily suggest that you test for control (negative feedback)
before you start developing models of what appears to be
cause-effect processes."

That's ALL we ask. Just test for control -- please! Without
such testing we have absolutely no way to evaluate the
success of any modeling efforts. We don't say cause-
effect models are wrong -- though we suspect that they are
most unlikely to be correct. But that is an empirical, not
a theoretical question. If all these cause-effect modelers
would just start doing the test for controlled variables, then
I, for one, would be happy -- whatever the results.

Regards   Rick

Date: Mon Mar 30, 1992 3:07 am PST
Subject: The Test, please!

From Greg Williams (920330)

>Rick Marken 920329 18:40

Right on, right on, right on!

Greg

Date: Mon Mar 30, 1992 3:38 am PST
Subject: Little Man model

From Greg Williams (920330-2)

>Bill Powers (920329.1400)

>We may as well go public with the arm-model discussions.

Regrettably, I won't be able to contribute much to the discussions for some
time. Over the next month, I need to finish up (the final, thankfully!)
CONTINUING THE CONVERSATION, get the galleys of Rick's book to him, do the
April HORTIDEAS, write half of the next Brooklyn Botanic Garden PLANTS &
GARDENS NEWS, and do the next CLOSED LOOP. Then there's a half-finished house
that needs my attention.... So, if you want to try to publish soon, I
recommend that you go ahead. I don't need to be co-author.

>I see what you mean by the outward curvature on upstrokes and downstrokes,
>although there is a lot more variation between subjects than was my
>impression from what you said. I wish we had the original data -- it's hard
>to get any quantitative measurements of what the two joints were doing from
>the figures. There's also a critical piece of missing data: in these plots
>of visual-motor behavior, the position of the eye isn't shown, and it's not
>mentioned whether the head bent forward as the arm descended!

Why not contact the experimenters for additional info?

>The most interesting problem is the speed of movement and the shape of the
>tangential velocity curve. When the traces of tangential velocity are
>normalized for duration and amplitude, they all have very nearly the same
>shape. This looks like gain control. In my model, there's no provision for
>controlling speed of movements. Perhaps, by putting gain control into the
>visual system, this effect can be reproduced. This would only be germane
>for the upper range of speeds, however.

What do plots of tangengential velocity vs. time look like for the little man now?

>If you introduce variations in the reference signals to the visual systems in
>my model, you can create any path you like.

But then how do you justify the particular path you choose? One way to avoid such criticism would be to show that if you picked a particular rule for setting the path of the moment-to-moment reference signals, then experimentally-seen paths would actually be predicted, regardless of where they started and stopped.

>The curvature problem is not so interesting. If the traces of fingertip
>movement were drawn with a line one centimeter wide, the difference from my
>model's behavior would look a lot smaller, particularly if you merged the
>data for all the subjects.

The different DIRECTIONS of curvature for the little man when moving "in" and "out" still bothers me. No great shakes, except that I suspect it would bother the reviewers, also.

>For my part, if the position of the fingertip stays within a centimeter of
>the average real fingertip throughout a movement I'd be satisfied. This model
>has only two levels in it, and no correction for distortions at all.

The reviewers probably would be a lot happier if you did whatever is needed to make the curvature predictions better. But of course you could go with what you've got and note the limitations. I just think the chances of getting it published would be improved if the little man (somehow) showed trajectory curves qualitatively similar to those in the data, in terms of curvature.

Again, I'm sorry I'm too busy right now to work with your model in detail.

Greg

Date: Mon Mar 30, 1992 7:19 am PST
Subject: Littleman arm model

[From Joe Lubin (920330.0900)]

Bill Powers (920329.1400) and Greg Williams

I've been spending lots of time working through motor control physiology and specifically your circuits. Now I am delighted to hear that you two have been having behind the net conversations about exactly that. Could you either post to the net what you may have saved from previous/recent relevant discussions or if this is too bulky could you send it to me email, although I'm sure others would benefit from a post.

Date: Mon Mar 30, 1992 8:58 am PST
Subject: Notes on arm model

[From Bill Powers (920330.0800)]

Greg Williams, Joe Lubin --

Notes on arm model.

Some solutions to the trajectory problem came together this morning. The model needs at least one more level, the transition level, which I now am convinced could also be called the "path" level. I've been rejecting the idea of path planning because those who use it in the literature have coupled it to the use of inverse kinematic calculations to produce torques as if the situation were open loop. But I now see that in order to create the kinds of tangential velocity profiles shown in the Atkeson and Hollerbach article, it's necessary to generate movements along a path in perceptual space (inside the model) with that velocity profile. Otherwise the distortions created in going from external Cartesian space to joint-angle space will (and do) create very strange trajectories of the fingertip.

The spinal systems in the model see to it that the actual joint angles follow reference signals for muscle stretch very closely. The actual kinematic properties of the arm are wiped out for all motions that take longer than about 0.15 sec to complete. So this level, which I suppose encompasses intensities (force) and sensations (stretch), is working satisfactorily.

The next level should control in configuration space. We could have a visual and a kinesthetic configuration space, but I'm using only a visual space. Including configuration space based on actual sensed joint angles and combining it with visual space is a project for the future.

In configuration space we have r, theta, and phi coordinates for the target and fingertip, the coordinates being angles and depth as reported by the visual perceptual system at this level.

By adding head angles (and later eye angles) into this information the visual configuration space can be defined relative to the body; I haven't done that yet. Kinesthetic configuration information, when that's added, can also be computed relative to the body, and then the perceptual functions in visual and kinesthetic modalities can be modified so they yield a common body-centered configuration space. All these are projects for the future.

What's been the problem in controlling the path of the fingertip for rapid movements is that when the fingertip moves in a straight line in objective space, it doesn't move in a straight line in perceptual space or kinesthetic space -- and vice versa. If the fingertip is to the right of the eyes and the target jumps horizontally to the left of them, a straight-line movement to the target would bring the fingertip closer to the eyes at first, then farther away again. As the fingertip gets closer to the eyes, the control systems push it farther away, so the path bows outward. But then, as the fingertip keeps moving to the left, it now is too far away relative to the target distance from the eyes, and the path bows inward again. The net result is that the fingertip doesn't naturally want to move in straight lines.

At first I thought this was a kinematic problem. Then I introduced the switch that allows the dynamics to be turned on and off, and found the same problem with dynamics off. It's simply a problem with the geometry of vision, coupled with the geometry of arm movement.

In order to bypass this problem for the time being, I used a "phantom target" that moved through intermediate positions from the initial to final positions. This target moves in Cartesian (objective) space in a straight line, because the x, y, and z coordinates of the target are set by a common parameter that goes from 0 to 1, scaled for each coordinate. If this phantom target moves slowly enough, the fingertip always stays on the target, and all the nonlinearities are eliminated by the feedback processes. And naturally, the fingertip moves in a straight line.

The Atkeson and Hollerbach article shows that straight-line motion is not required (at least when subjects aren't told to follow any particular path). The "default" paths are somewhat curved, depending on the placement of the end-points. Rough measurements of their plots show that the radial distance from the shoulder simply changes from the initial radius to the required final radius in a smooth curve, as does the shoulder-to-fingertip angle. While the data aren't enough to verify this, it would seem that the movement of the fingertip in radius and vertical angle is controlled parametrically -- that is, as if the finger reference position is being varied simultaneously in these two dimensions by a change in a common parameter. Velocity profiles indicate that this supposed parameter traverses its range in a way that accelerates to the midpoint, then decelerates to the final point. The authors indicate that this is a "minimum-jerk" movement, meaning that the rate of change of acceleration is smooth.

All this is very suggestive. To me, it suggests that we need one more level, a transition level, to control the parameter that moves the fingertip reference position from the initial position to the target position, along a straight line in perceptual space. This will be a curved line in objective space, at least for some directions of movement (radial changes will be along straight lines). The transition control system can easily be set up to produce the same velocity profile for all movements and speeds of movement; this will naturally come out of the feedback dynamics. I can see now that we will be able to make a prediction: the velocity profile, normalized as in the article, will also be the same for different spatial separations of initial and final positions, at least for rapid movements.

As the model is set up, a geodesic movement in perceptual space will be transformed into an eye-centered movement in objective space. According to the Atkeson and Hollerbach data, the actual movement seems to occur in a space that compromises between shoulder-centered and eye-centered. This is very hard to judge, though. In the end I think we will have some sort of body-centered space with both visual and kinesthetic perceptual functions being modified to reconcile them with a common space. Then a geodesic movement will be straight in this common space, but still probably curved in objective space. True straight-line movement in objective space will have to wait for the relationship level to be built, I think. Something will have to determine that a particular curved path in configuration space has especially nice properties when you follow it. Or maybe this space eventually adapts to achieve some minimum-energy-of-action state, and distorts until straight lines are really straight. All this ought to be great fun to work out, but I hope my poor muddled brain isn't called on to do it all.

Of course once we can produce geodesic path control, we can produce any other kind. All that has to be done is lay out a different path in perceptual space, and let the transition-control system run the reference signal settings for the fingertip along that path. This is starting to sound pretty real, isn't it? Now we can have path-planning that simply amounts to tracing out a path in imagination and then making the fingertip position reference signals trace it in r, theta, and phi. Everything will be under tight feedback control all of the time.

I think this addition will make the model more acceptable to conventional modelers, because it does allow for trajectory planning. It eliminates the need for those awful inverse kinematics computations, which means that a simulation or hardware model actually has a chance of working (for indefinite periods of time). I don't think that anyone has actually tried running a model based on the inverse kinematics approach, at least not for any extended period of time. Those integrations would get out of whack pretty fast!

I've had vague ideas that the transition level is involved in path control, even in BCP (p. 133), but this concept seems closer to realization now than it's ever been.

Joe Lubin, I hope you continue to be interested and can find more students willing to extend this model-building. I have such mathematical limitations that I really can't do all of the things that I can envision. When I get into coordinate transformations and such I feel that I'm whipping a reluctant old horse through a thick fog. It's really getting to be time for some young sharp brains to get into this act.

Greg, how much of this do you think we actually need to get working before we submit a paper on this model? I'll try to get some kind of parametric path control into it, but I do think it's time to get into print. The Science referees on that letter in effect challenged me to put up or shut up, so I think there's a good chance of publication there.

Joe, How about some details on what your student has accomplished so far?

Best   Bill P.

Date: Mon Mar 30, 1992 10:06 am PST
Subject: Re: Category, configuration, sequence control

[Martin Taylor 920330 11:15]
(Bill Powers 920329.1400)

No time for proper reply. Panic is setting in on paper deadlines!

>Thinking generally about your problem of sequence control, I've had a
>thought that may be useful. It's similar to a thought I've had about
>configurations, so I'll start with that.

I won't quote all your stuff, but yes, I think we are on a useful track. In an HCI paper I am panicking about (not the big one), I describe the ECS, and then suggest that a group of same-level ECSs could be looked at together as if they had a vector percept, reference, and error in one box. I called the group a Structured Control System (SCS). If such a set of ECSs were linked so that their references formed a sequence, the same concept would presumably apply over a stretch of time, which would give the same sort of result as your sequence-level ECS, but would make more conceptual sense to me. I think this is converging with what you are saying.

If I can persuade myself, I will have very little to say here until June, but I'll be keeping the postings. I'll actually be in town this week and April 12-26 or thereabouts.

Martin

Date: Mon Mar 30, 1992 11:09 am PST
Subject: S-R vs CT; Entropy & such

[From Bill Powers (920330.0930)]

Greg Williams (920329) --

>Are you claiming that there is actual data showing that (a) in the case
>of short-duration air puffs, the body orientation changes before the >air
puff is over, and/or that (b) if the body orientation does change >before
the (say, somewhat longer-duration) air puff is over, the >movement of the
cockroach is actually influenced by the change in >alignment between the
body and the continuing air puff?

I'm making a deduction. Beer said that the turning movement was completed in 60 milliseconds. A movement of air caused by a large approaching clodhopper would without doubt last a lot longer than 0.06 sec. So it's dead cert that the movement of air around the sensing hairs is affected (strongly) by the turning movement before the movement of air has ceased, and even before it has peaked. Actually I would expect such a surge of air to propagate outward for several seconds, the cockroach not only turning while the air movement is still going on, but moving downwind a considerable distance, creating local effects that probably exceed the amplitude of the original air puff. There can be no question that the behavior of the cockroach strongly influences and even cancels the effects the air puff would have if the cockroach stayed in one place. So this is a negative feedback situation, however the cockroach is organized to behave in it.

Because nobody has done anything like a quantitative experiment with this escape response (I'll bet), nobody knows whether the cockroach is acting as a control system or just goes through a fixed repertoire of actions. The only reasons to reject the closed-loop hypothesis, so far, go like this: feedback is too slow; evolution just says "get out of this place;" the response (escaping) doesn't affect the stimulus (stamping your foot). Are those reasons good enough for you?

>There ARE a lot of cases where g(r) is demonstrably nonzero. But it is
>quite a leap to your general claim. I wish you could see that the worth
>of HCPT modelling does not hang on the truth of that generality.

No, the giant leap is the proposal that there are any cases in which behavior does not immediately influence the stimuli that actually "caused" it. The real worth of the HPCT (sic) model is in the proposition that strong negative feedback loops are the foundation of all organized behavior. If it were true that only a few systems here and there showed this interesting property, then PCT would be just a curiosity and there would be no reason to re-examine any S-R formulation that seems to work (no matter how poorly). And HPCT would be nonsense, because you can't have a hierarchy of control that includes a lot of S-R systems that behave without regard to the rest of the system. They would just be disturbances, and the closed-loop systems would resist their actions.

None of this means that we have to close our minds against the possibility that some output is caused by some input that isn't in a feedback relation to that output. Sure, it can happen. But so what? Are we to give up looking for closed loops just because we can see a situation as open-loop? Like the foot-stamping causing the escape response with no feedback effect on the foot-stamping? If there really are straight-through responses, the methodology of control theory will find them. Feedback-controlled variables are not hypothetical; they're testable. If those apparent straight-through responses result merely from taking too anthropocentric and general a view of the situation, S-R theory and its methods will NEVER lead to discovery of the truth. There is no point in assuming S-R connections, and there is every reason to assume closed-loop connections until the data prove otherwise. This is true even in the most solid-appearing examples of open-loop behavior.

This has very little to do with humility on my part (perhaps for good reason, like hubris). It has a great deal with seeing how poorly behavior has been observed in the past, and with what giant leaps of faith the S-R model has been defended.

Martin Taylor (920329.1330) --

>I'm not making an explanatory model of anything when I refer to the maintenance of entropy. I'm describing negative feedback. >Environmental disturbance increases the entropy of the system, and >negative feedback reduces it. In a stable controlled system the >entropy remains stable despite disturbance.

If entropy is controlled, then it's sensed, compared with a reference entropy, and maintained at that reference level by variations in action. I don't buy it. I don't even think that entropy is systematically affected by negative feedback control of any perceptual variable. If you hold your arm out straight, the entropy of your system increases more and more rapidly as fatigue sets in; you then have to rest in order to absorb or free from storage some negative entropy to make up for the drain.

Entropy, like information or probability, is a calculation by an observer, not a system variable. It changes as dQ/Q, if I remember right. This calculation can be applied to any measure of anything. It's controlled only if the behaving system specifically senses it (calculates this function of its inputs) and acts to keep it at some preferred level. An observer's calculations of the entropy of some other system have no necessary relevance to what makes that system work. If the calculations prove consistent, then all that has been proven is that there's a consistent side-effect of the action of the system. This does not make the side-effect either causal or explanatory of the system's operation.

I just don't buy explanations of organized systems that relay on abstract measures. Organized systems result from real interactions among real variables, not from abstract characterizations of those interactions.

>Yes, I mean that disturbances, within limits, will be resisted in such >a way that the disturbed structure returns to its undisturbed state.

This is still too qualitative a statement. Does the disturbance cause any signficant departure from the undisturbed state? Can you cite an example of what you're talking about? Even a pendulum would fit your statement, but it's certainly not a negative feedback control system. All the restoring energy, as Gary Cziko mentioned, comes from the disturbance, in the pendulum.

>I maintain that the distinction you are looking for is not the high->gain negative feedback, but the reference-signal + comparator >structure.

I'm not looking for a general abstract distinction, but for a description of the way organisms work. High gain negative feedback (with or without a variable reference signal, with or without an explicit comparator) creates the kind of behavior we see in living systems at all levels from biochemistry to control of system concepts.

>What I hope to do, probably starting when I get back from my trip in >June, is to come to a resolution of the problems of the thermodynamics >of the control system hierarchy.

Fine.

>... a subjective lag of zero is not the same as a loop delay of zero.
>The latter refers to how long it takes the perceiving system to acquire
>information about the effect of action in the presence of all sorts of
>environmental disturbances and disturbances induced by competing ECSs.

Agreed. There is, however, one objective criterion that must be met: whatever the actual loop delay, something in the loop must insert averaging of such an amount that the system would behave no differently if the actual lag were zero. Otherwise the loop can't be stable.

>Why do you say there are no "internal" loops? Do you mean that there >are no loops within a single ECS?

Yes -- no behavioral loops. Any individual function (input, comparison, output) could contain feedback loops as part of the computational process, but the ECS as a whole has no internal loops.

>Do you mean that imagination doesn't happen?

No. Imagination results from outputs of an ECS that enter an imagination connection external to the ECS and arrive back at its input as if a real perception were being received. This connection is external to the ECS as I define it: input, comparison, output.

>Or that kinaesthesis doesn't happen?

Kinaesthesis happens. The output of the spinal ECS affects muscles in its environment. Those muscles have physical effects on local tissues, where effects of independent disturbances also can occur. The result is excitation of sensory nerves at the inputs of the ECS. I place the boundary of a higher organism's behavioral systems so that everything outside the nervous system is in the environment of the behaving system.

>Or are you making the assertion, that I expect to take issue with >later, that there are no cross-connections within a level of the >hierarchy?

I make no use of such cross connections but I don't deny their existence. Some connections that appear to be cross connections are really composite control systems -- for example, the connections that make a stretch response on one side of a joint relax the muscle on the other side. The opposing muscles are part of a single control system. The only true cross connections, I think, in the terms you mean them, would be between systems at the same level which are otherwise independent of each other. Such connections exist but I have only a vague idea of what they accomplish.

>The chaotic behaviour (or, more probably, near critical behaviour) is
>required for the perceptual functions, and it is to make that behaviour
>useful that we have categories. I think that's what Freeman is saying,
>as well. But even if he isn't, I am.

I think even perceptual functions can become organized through random, not chaotic, change. Freeman's concept of chaos in perception is that there are modes of operation of a whole chunk of the brain, like the olfactory lobe, and that perception is the existence in this chunk of some mode of large-scale distributed oscillation. In the absence of inputs, this large system goes into a chaotic mode, which can then settle into new basins of attraction when novel input combinations occur, thus leading to a new mode of perception.

While Freeman's waving his arms and conducting this music, the composition sounds inspiring and beautifully constructed and the words to the song are very plausible. But after the conclusion, the questions rush in. What good does it to to have the whole olfactory system oscillating? How does that lead to following a scent to its source? What is it that discriminates one mode of oscillation from another, saying "That's chocolate" and "That's perfume"? Where's the recognizer? Whqt does all this have to do with perception? Somehow Freeman has managed to show that a large-scale mode of oscillation can depend on inputs, but he still hasn't answered the question of how we distinguish one input from another. I think he does some very intensive interpreting, assuming, and even imagining when he reports the "facts."

But I have to admit that there's a kernel of an idea there that's attractive. One wants to believe that somehow our perceptual worlds are truly basically alike. If there are mathematical laws that say our perceptions naturally settle into certain preferred forms, rather than just being scattered at random among the possibilities, then there are grounds for thinking our perceptual worlds are pretty similar. I guess I just have to wait on the sidelines for further (and more convincing) developments.

Best to all   Bill P.

Date: Mon Mar 30, 1992 12:39 pm PST
Subject: Re: Open loops, closed loops and HCI

[Martin Taylor 920330 11:45]
(Rick Marken 920325 13:00)

Sorry for the long delay in responding. This thread is work, and I would like to get it right, in June. But I want to respond to one point Rick made:

>Martin mentioned two in his comments [...]
>2) "error reduction" is one of the big concerns in the field (others
>include "efficiency", "usability", "safety"); obviously, HCI people
>typically use the term "error" to describe descrepencies between what
>they think a result should be and the result being produced by the operator.
>Thus, the "errors" are experienced by the HCI engineer, not necessarily
>by the operator

This is a very astute point, and brings out part of the design problem. The HCI engineer presumably has a reference for what the operator should do (action, not behaviour) in particular circumstances. Failure of the operator to act that way is "error" that can be "corrected" only in the fashionable loop of prototype, test, modify, test. But the failure is probably of the HCI engineer to realize that the operator is behaving, not acting. What percepts will the operator be comparing with what references at this stage in the ongoing interaction? What is the operator "wanting to do?"

Part of the problem is what we call "situation awareness" on the part of the operator. How can the HCI engineer arrange that the computer display contains appropriate (and appropriately few) alerting signals--patterns that "automagically" bring the operator to begin to control percepts that were up to that point being passively (and perhaps unconsciously) accepted?

"Errors," from the viewpoint of the HCI engineer, can come from at least two disparate causes: the operator does not control the percepts the engineer would wish, or the operator does attempt control but is improperly organized to make that control effective. In the first case, alerting (focusing) signals are required, and in the second, reorganization (of the operator).

Neither alerting nor reorganization have been strong threads in this group. Both, it seems to me, depend on the degrees of freedom arguments that I have slowly been trying to develop. I wish I had more time right now to devote to it, but matters press.

Martin

Date: Mon Mar 30, 1992 12:42 pm PST
Subject: Re: Notes on arm model

[Martin Taylor 920330 13:20]
(Bill Powers 920330.0800)

>What's been the problem in controlling the path of the fingertip for
>rapid movements is that when the fingertip moves in a straight line in
>objective space, it doesn't move in a straight line in perceptual space
>or kinesthetic space -- and vice versa. If the fingertip is to the right
>of the eyes and the target jumps horizontally to the left of them, a
>straight-line movement to the target would bring the fingertip closer to
>the eyes at first, then farther away again. As the fingertip gets closer
>to the eyes, the control systems push it farther away, so the path bows
>outward. But then, as the fingertip keeps moving to the left, it now is
>too far away relative to the target distance from the eyes, and the path
>bows inward again. The net result is that the fingertip doesn't
>naturally want to move in straight lines.

Actually, there isn't a "perceptual" problem of this kind. The idea that there is stems from an insufficient adherence to the ideas of PCT. As JGTaylor, and Kohler before him, showed, very strong distortions of the visual space by trick spectacles initally result in corresponding distortions of perceptual space, but PROVIDED that the wearer was able to control those perceptions, the perceptual space rapidly regained its Euclidean character. Only with perceptual control did this happen. Passive exposure to the distored space, even with movement (e.g. being wheeled around in a chair) has little effect in removing the perceptual distortion.

For people in a natural world, all the distortions such as Bill mentions will have been eliminated in favour of Euclidean perception. Maybe this relates to the idea of path planning?

Martin

Date: Mon Mar 30, 1992 1:01 pm PST
Subject: Re: S-R vs CT; Entropy & such

[Martin Taylor 920330 13:30]
(Bill Powers 920330.0930

Bill, I think you answered your own criticism of my comments on entropy, so I probably shouldn't comment. But anyway, to clarify...I make no suggestion that entropy is controlled in the sense that a percept is controlled. A pendulum does indeed reduce its entropy, by exporting the energy of the disturbance. Reduction of entropy in a strong energy flow is a sign, a demonstration, that negative feedback is occurring, and stabilization of entropy below a maximum value is a sign that the negative feedback system is near its attractor, which in a control system would be described by its reference signal.

OK? We agree? I think I restated what I said before, in words closer to your statement.

As for whether a disturbance causes a "significant" departure from the undisturbed state--significance is in the eye of the beholder. If one disturbs a dynamic structure to a point near the edge of its attractor basin, and it returns (i.e. negative feedback), has the disturbance been "significant?"

I think that's a non-question. If I break your arm, and your finger no longer tracks the target, has that been a "significant" disturbance? If I just bump your arm slightly, and you quickly regain track, was that disturbance "significant?"

I did cite an example--the vortex in a strong fluid drain. As for the pendulum, it's not in a strong energy flow. A one-shot deposit of energy is dumped into it, and is dissipated. Uninteresting.

>I'm not looking for a general abstract distinction, but for a description
>of the way organisms work. High gain negative feedback (with or without a
>variable reference signal, with or without an explicit comparator) creates
>the kind of behavior we see in living systems at all levels from
>biochemistry to control of system concepts.

Yep, but not without reference signals, I think. As to whether comparators are explicit, I'm not sure what that would mean except in a simulation.

>There is, however, one objective criterion that must be met:
>whatever the actual loop delay, something in the loop must insert averaging
>of such an amount that the system would behave no differently if the actual
>lag were zero. Otherwise the loop can't be stable.

That's too strict a criterion. There are lots of ways to stabilize systems with delay. It's bandwidth and phase response that count in a linear system, and who knows what in particular kinds of non-linear systems.

>>The chaotic behaviour (or, more probably, near critical behaviour) is
>>required for the perceptual functions, and it is to make that behaviour
>>useful that we have categories. I think that's what Freeman is saying,
>>as well. But even if he isn't, I am.

>I think even perceptual functions can become organized through random, not
>chaotic, change.

Once again, no-one is (yet) suggesting chaos for reorganization. The pending claim is that it is required for the kind of rapid response at high levels that you were talking about the other day. The other claim, which is aimed much more at the AI folks than at CSG, it that it is essential, together with catastrophe functions, if one is going to perform logical or categorical operations. Coming at things from the CSG side, I think one can ignore that facet.

Martin

Date: Mon Mar 30, 1992 4:43 pm PST
Subject: Taylor stuff

[From Bill Powers (920330.1600)]

Martin Taylor (920330.1320)--

I promise not to open any new subjects -- I know you're trying to get unhooked so you can do all those things you REALLY HAVE TO DO now. Don't even reply to this ---

Re: perceptual distortion in arm model.

>Actually, there isn't a "perceptual" problem of this kind. The idea >that there is stems from an insufficient adherence to the ideas of PCT.

Oh, yes there is -- in MY arm model. I don't have any reorganizing abilities in the model, or any way of mapping from one space to another (except Cartesian to r-theta-phi, which is trivial). I've been sorely tempted to cheat and say "Oh, well, the brain will adapt itself and provide x,y,z signals anyway, so why not just use the external coordinates for the perceptual signals?" Actually it's much more instructive to go through all this and see just what you can get away with without any adaptation at all -- and what jumps up and bites you. In principle, the perceptual space doesn't have to be modified: all that has to change is the set of points that's called "a straight line." If the reference signals trace out the correct set of points, the result will be a straight line in external space. Right now I'm tinkering with the innards of the program, trying to find out what's working well enough to leave alone and what has to be changed to introduce the higher levels. Hand me that wrench, will you?

I guess we agree on entropy. When you get back we'll see. An entropy measure can indicate the success of the control system. Sort of like error as a fraction of the total range...?

Significant disturbances:

>As for whether a disturbance causes a "significant" departure from the
>undisturbed state--significance is in the eye of the beholder.

The rule of thumb I use is based on the total range of the controlled variable, which is the range of the reference signal. A sort of nominal control system can counteract any disturbance large enough (when unopposed except by the passive dynamics of the controlled variable) to drive the variable to its limits (if it has enough output capacity to cancel such a disturbance). A "large" disturbance of a pendulum would be one that pushes it 90 degrees from vertical (less force is needed to push it further). Control is "good" if that size of disturbance is kept from varying the controlled variable more than 10 percent of the total range. That is, the same force that would push the pendulum 90 degrees without control now moves it only 9 degrees. This is obviously an arbitrary measure. It's conservative, though, in that most compentent control systems I come across in behavior can do better than this.

>If one disturbs a dynamic structure to a point near the edge of its
>attractor basin, and it returns (i.e. negative feedback), has the
>disturbance been "significant?"

By my definition, yes. One of the natural ways to define a "large" disturbance is in terms of the range of the affected variable, which in complex systems is usually finite. If a measure of the state of the system is driven all the way to a limit, the disturbance is large. Also, the resistance of the system to disturbances is weak.

This is a somewhat tricky question because "disturbance" is ambiguous. It can mean either the cause or the effect. In CT we almost always mean the cause, the independent variable that INFLUENCES the controlled variable. If you grab the controlled variable and force a change, this breaks the loop. So I translate your disturbance to mean "a change in an independent variable that results in the dynamic structure moving to the edge of its attractor basin." This means that a control system could vary a second independent variable that has the opposite effect, and thus maintain the dynamic structure at the center of its basin even in the presence of the other disturbing influence. It could also keep the dynamic structure at any arbitary distance from the center of its basin, in the presence of abitrary disturbing influences.

This is the very question that Normal Packard, at the U of IL, copped out on many a moon ago. I had asked him what the effect on one of his dynamic systems would be if a constant disturbance were applied. He said he would answer "after the holidays" and never did.

Re: reference signals

ME:
>>   ...High gain negative feedback (with or without a
>>variable reference signal, with or without an explicit comparator)
>>creates the kind of behavior we see in living systems at all levels
>>from biochemistry to control of system concepts.

YOU:
>Yep, but not without reference signals, I think. As to whether
>comparators are explicit, I'm not sure what that would mean except in a
>simulation.

All you need for a control system, beside external feedback and dynamic stabilization, is a system in which output = -K(input), assuming an external feedback connection with a positive constant. This system doesn't need a comparator and it has no provision for a reference signal. It will keep its input near zero. It is, of course, equivalent to a control system with a comparator and a reference signal set permanently to zero. It couldn't be a subsystem in a hierarchy of control, however. But it's a control system, if the loop gain is high and negative.

Reference signals can be added into perceptual functions or output functions (the latter is how some brainstem systems seem to work). There's no need for a separate circuit that does the subtracting. It's just easier to understand the system if there is one.

>>something in the loop must insert averaging of such an amount that the
>>system would behave no differently if the actual lag were zero.
>>Otherwise the loop can't be stable.

>That's too strict a criterion. There are lots of ways to stabilize
>systems with delay. It's bandwidth and phase response that count in a
>linear system, and who knows what in particular kinds of non-linear
>systems.

My criterion is equivalent to yours, and applies in the nonlinear case as well. If the Laplace transform has a delay of tau seconds, it contains a term exp(-tau*s). The system is stabilized by a single leaky integrator with a time constant of tau. The result is that the system behaves like another system with no delay and no leaky integrator. Probably I should have mentioned that it's the low-frequency behavior that's the same -- of course with no delay and no filtering, a system could respond infinitely fast. I just meant that if the system is stable, its delay can be ignored in computing its behavior over the frequency range of the stabilized system (as long as you don't get too close to the limits of performance).

Re: chaos

>Once again, no-one is (yet) suggesting chaos for reorganization. The
>pending claim is that it is required for the kind of rapid response at
>high levels that you were talking about the other day.

You'll have to develop that thought a bit further before I see the connection.

Go away. Have fun in Paris.

Best, Bill P.

Date: Mon Mar 30, 1992 5:27 pm PST
Subject: Closed Loop/Arm/Roach Escape

From Greg Williams (920330)

EPISTEMOLOGY (and another topic, if there's room, but it's looking like there won't be) will comprise the subject matter in the next CLOSED LOOP, due out this month, by yours truly's unilateral decree. Those who don't like unilaterality (and even those who do) are welcome to suggest possible subjects for the July and later issues.

>Bill Powers (920330.0800)

>Some solutions to the trajectory problem came together this morning. The
>model needs at least one more level, the transition level, which I now
>am convinced could also be called the "path" level.

I suspect that you can get it to work, even though it isn't as elegant as pure end-point control.

>All this is very suggestive. To me, it suggests that we need one more
>level, a transition level, to control the parameter that moves the
>fingertip reference position from the initial position to the target
>position, along a straight line in perceptual space.

Sounds like a good rule to begin with. Go for it!

>I can see now that we will be able to make a prediction: the velocity profile,
>normalized as in the article, will also be the same for different spatial
>separations of initial and final positions, at least for rapid movements.

When confirmed, that will get some attention!

>I think this addition will make the model more acceptable to
>conventional modelers, because it does allow for trajectory planning.

Right on!

>Greg, how much of this do you think we actually need to get working
>before we submit a paper on this model?

I don't know about the "we" business -- I'm just the guy who found the equations of motion, remember? -- but the choice seems pretty clear: end-point control vs. path control, and I still don't think the referees are going to like the shapes of the end-point-control trajectories. On the other hand, if you were up front about the "limitations" of end-point control and talk about how well path control would work, maybe they'd accept the model as it stands now. In suppose it depends greatly on the backgrounds of the referees. Don't wait on my account.

>Bill Powers (920330.0930)

Greg Williams (920329) --

>Are you claiming that there is actual data showing that (a) in the case
>of short-duration air puffs, the body orientation changes before the >air puff is over, and/or that (b) if the body orientation does change >before the (say, somewhat longer-duration) air puff is over, the >movement of the cockroach is actually influenced by the change in >alignment between the body and the continuing air puff?

>I'm making a deduction. Beer said that the turning movement was completed
>in 60 milliseconds. A movement of air caused by a large approaching
>clodhopper would without doubt last a lot longer than 0.06 sec. So it's
>dead cert that the movement of air around the sensing hairs is affected
>(strongly) by the turning movement before the movement of air has ceased,
>and even before it has peaked.

You have deduced that cockroach movement will affect the influence of the air-puff on the hairs. But is the loop continuously closed? You have NOT deduced that the movement of the cockroach is influenced by the changes in influence of the air-puff on the hairs during the movement. That is the part of the feedback loop (if it is a loop) inside the organism.

>There can be no question that the behavior of the cockroach strongly
>influences and even cancels the effects the air puff would have if the
>cockroach stayed in one place. So this is a negative feedback situation,
>however the cockroach is organized to behave in it.

Yes, we KNOW that there is an overall feedback loop. But is a particular cockroach escape action the result of an internal pre-calibrated (by evolution and possibly learning) chain mechanism, or a continuous control-of-perception? That is the $64 question. As I posted to Rick recently, in the case of a chain mechanism, all I think is compromised with respect to PCT ideas is the notion of control always being CONTINUOUS.

>Because nobody has done anything like a quantitative experiment with this
>escape response (I'll bet), nobody knows whether the cockroach is acting as
>a control system or just goes through a fixed repertoire of actions.

I'll be in Lexington tomorrow to do research for HortIdeas; maybe I'll have time to look for evidence one way or another in Camhi's book and elsewhere.

>The only reasons to reject the closed-loop hypothesis, so far, go like this:
>feedback is too slow; evolution just says "get out of this place;" the

>response (escaping) doesn't affect the stimulus (stamping your foot). Are
>those reasons good enough for you?

There is a potentially good reason why evolution might "just" say "get away": the neural overhead associated with PCT-circuitry might be more "costly" than than associated with chain-circuitry. But I'd need to see the particulars before claiming that any of these reasons, construed generally, is "good enough."

>No, the giant leap is the proposal that there are any cases in which
>behavior does not immediately influence the stimuli that actually "caused"
>it.

That WOULD be a giant leap. The less-giant leap which actually seems to be the one taken by several nonPCTers is that actions (I think that is what you meant, not PCT-type behavior, which would make the "giant leap" simply impossible by definition!) DO immediately influence the "stimuli," but often that influence doesn't affect the trajectory (taken generally) of the action.

>The real worth of the HPCT (sic) model is in the proposition that
>strong negative feedback loops are the foundation of all organized
>behavior.

Fine by me. The "hangup," as I see it, is the insistence that control be continuous.

>you can't have a hierarchy of control that includes a lot of S-R systems that
>behave without regard to the rest of the system. They would just be
>disturbances, and the closed-loop systems would resist their actions.

If a pre-calibrated chain mechanism contributed to higher-level control-of-perception, it wouldn't be resisted. And if there were no reference levels at the level of the chain mechanism, it wouldn't act like a disturbance.

Are we to give up looking for closed loops just because we can see a situation as open-loop?

No, we should be looking for higher-level loops. And setting aside the claim that all control is continuous.

Greg

Date: Mon Mar 30, 1992 5:47 pm PST
Subject: mpd

To: Mark Olson and others interested in MPD
From: David Goldstein
Subject: MPD & HPCT
Date: 03/30/92

Yesterday, when I answered Mark Olson's post of 03-27-92, I forgot to address one question which he posed.

Question: Is there one thing being fulfilled by these multiple self-concepts or not? I think so. In my patient, there is an alter which we have variously called "the controller, or the keeper of the basket, or the overseer" who knows everything about the other alters and takes some responsibility for their creation. Within a "normal" person, I have speculated that there is "an observer" who selects which one(s) of the self-images will be operant in a given circumstance.

Best regard,
David Goldstein
internet: goldstein@saturn.glassboro.edu

Date: Tue Mar 31, 1992 2:58 am PST
Subject: Re: S-R theory its own self

[From Oded Maler (920331)]

>[From Bill Powers (920328.1030)]

Your demonstration of the difference between S-R and Control is a mixture of two orthogonal dimensions. Now I'm almost convinced that (in principle) S-R and control are identical phenomena. Some of the arguments against S-R are not against the principle but against the practive of S-R psychologists. i'll try to clarify:

The main argument in your post was a critique of "anthropormpization" of perception and action, that is, characterization of events from the point of view of an external observer rather than from the perspective of the behaving organism itself. I agree with you completely on this topic, and if behavioral psychologists usually ignore this fact, this discipline is not worth much. Recall however, that if you climb up above the most basic sensation, even your models cannot escape from this problem as long as the categorization problem is unsolved. When you say that someone controls for being far away from a "dog", you can say *in principle* that "dog" can be somehow represented by internal perceptual coordinates of that someone, by you'll never really have such internal descriptions.

Even in this quote one can find a shadow of anthropormpization:

>This modified the sensor signal, and the response, which
>had been aimed at one final state, is now aimed at a slightly different
>final state.

The second argument was the "essential" difference between

r=f(s,s*) and r=f(g(r),s,s*)

If you just add time indexes (continuous, of course..) to r and s, and write

r[t]=f(s[t],s*)

then, by assuming that s[t] is somehow influenced by r[t'] for all t'<t, you can get an S-R formulation. If the organism can tell the difference between changes in s caused from the outside and the same changes in s

caused by its own actions, it just means that you should replace s by a more refined perceptual signal that can make these distinctions.

--Oded

Date: Tue Mar 31, 1992 6:03 am PST
Subject: Re: S-R vs CT; Entropy & such

[From Marcos Rodrigues]

(Bill Powers Mon 30 Mar 1992 10:39:11 -0700):-

>None of this means that we have to close our minds against the possibility
>that some output is caused by some input that isn't in a feedback relation
>to that output. Sure, it can happen. But so what? Are we to give up looking
>for closed loops just because we can see a situation as open-loop? Like the
>foot-stamping causing the escape response with no feedback effect on the
>foot-stamping?

I think the foot-stamping situation is in closed loop. The cockroach action changes its perception of the environment, surely. My understanding of your theory is that an action modifying _physically_ the environment is not required,

provided that our perceptions of the environment change in direction of correcting the error.
The cockroach overdriving response could simply be due to a time-dependent higher level loop (active say, for 500 ms) which drives the legs as fast as possible, while the direction of movement is dependent on which sensors are disturbed by the puff of air, vibration, etc. I can't see feedforward or open loop in the escape response.

Best wishes,
Marcos.
mar@uk.ac.aber

Date: Tue Mar 31, 1992 2:00 pm PST
Subject: Arm; S-R

[From Bill Powers (920331.1030)]

Greg Williams (920320) --

Good topic for Closed Loop. I'm continually amazed at all the things you do, and so well. You would make a great role model for any 4 people.

Arm model:

I'm putting in a second kinesthetic level of control. This level perceives the joint angles in such a way as to make radial movements relatively independent of movements in elevation. Also I'm putting "size constancy" into the visual perceptions, by making the x and y perceptions proportional to the distance perception: this will tend to make the gain of visual control loops constant for objects at various distances, and may help with the end-point control mode. This second level will provide a place to add the "artificial cerebellum" which you've already seen, to adapt the stabilization to various load conditions. I don't think I'll actually add the A.C. in this version.

Cockroach, SR/Control:

>You have deduced that cockroach movement will affect the influence of >the air- puff on the hairs. But is the loop continuously closed? You >have NOT deduced that the movement of the cockroach is influenced by >the changes in influence of the air-puff on the hairs during the >movement. That is the part of the feedback loop (if it is a loop) >inside the organism.

You're right that verifying the external feedback doesn't prove that the organism acts continuously. But it's possible to set up a general model that includes the feedback connection, and deduce the actual input-output transfer function (guessing, of course at the exact effect of movements on the wind-sensing hairs). The model of the cockroach's input-output function that makes the model behave the most like the cockroach, based on the ACTUAL input, will show whether the response is modified by the effect of behavior on the input while the response is occurring.

This is what makes the negative feedback control model more general than the SR model. If all the components of the model are left general, with parameters to be determined by the data, then in principle the parameters that result will show whether SR theory is adequate. For example, if the output-to-input connection is in fact missing, then the best-fit model should evaluate the feedback connection g(r) as zero. That's in principle, of course -- in practice that might be difficult.

There are some behavioral experiments that would make matters clearer. For example, if the puff of air is modulated to prevent the movement of the body from having the normal effect on relative air speed, then the final direction of "escaping" should change. Or if an obstacle is placed so that the escape response would make the cockroach collide with it, we could tell whether the direction of escaping would be modified to avoid the obstacle. We could put the cockroach into a passageway too narrow to turn around in and give it a puff in the face. We could glue a thread to the cockroach and just as the puff occurs, give it a tug that aids or opposes the initial turn. We could glue the cockroach's body to a support and watch what the feet do as the puff of air arrives from various directions. We could use a strong puff of air from the side applied so it aids or opposes the turning of the body, and see if the cockroach compensates for the spin induced by the puff of air. We could put a wall downwind from the cockroach with a hole through it and see if the cockroach modifies its direction of running to go through the hole wherever the hole is placed. We could put the cockroach on a slick Teflon surface and see if it compensates for the slipping of its feet.

I think that if anyone seriously wanted to test the choice between an open-loop and a closed-loop escape response, it would not be hard to think of experiments that would settle the question. Any takers out there among the bug people?

>There is a potentially good reason why evolution might "just" say "get
>away": the neural overhead associated with PCT-circuitry might be more

>"costly" than than associated with chain-circuitry.

What overhead? It's no simple matter to arrange for an open-loop behavior that turns the body by an amount that depends on the direction of the threat, then institutes running movements with the same legs. You have to have wind-sensing hairs all around the body. Then you have to connect them to a central computer so that if the wind is from angle A relative to the direction of the body, the turning-pattern generator starts lifting and moving the legs in just right sequences and by just the right amounts to produce a turn of 180 - A degrees. Then when this sequence of leg movements is finished, the same computer has to produce signals that cause fast forward locomotion for some period of time. On a bumpy surface.

All of this would be FAR easier to do with feedback control. No accurate computations would be needed. The only thing that's less complex about the "just get away from there" explanation is the thought process of the explainer. The extra overhead is in the S-R model, not the feedback control model.

>The less-giant leap which actually seems to be the one taken by several
>nonPCTers is that actions (I think that is what you meant, not PCT-type
>behavior, which would make the "giant leap" simply impossible by
>definition!) DO immediately influence the "stimuli," but often that
>influence doesn't affect the trajectory (taken generally) of the >action.

"Often" is a pretty giant leap. "Sometimes" I might buy as a possibility. But considering the defects of observation behind all S-R models, I would rather approach an unknown situation assuming control by feedback, and let the experiment show this is wrong. As Mary observed while reading this morning's posts, you always approach data with some model in mind. The S-R model is a pretty poor one to assume.

Mary also pointed out that human beings sometimes seek out situations in which there are trajectories uncontrollable by the actor once they are started: golf, bowling, baseball, basketball, archery, and skeet-shooting would be examples. All that can be controlled during the action is the delivery or initial aim in these cases; long-term control can be achieved only by repeated tries with small adjustments of reference signals between them. Controlling in this way is very hard and slow, which seems to be the appeal -- achieving control is very difficult so that when it's achieved, the person feels that something worthwhile has been accomplished. If control were easy there wouldn't be any element of a game or competition in it. If human beings could make every drive in golf a hole-in-one, nobody would play golf. On the other hand, if survival depended on playing any of these games perfectly, the games still wouldn't be played, but for a different reason. There would be nobody to play them.

>If a pre-calibrated chain mechanism contributed to higher-level >control-
of- perception, it wouldn't be resisted. And if there were no >reference
levels at the level of the chain mechanism, it wouldn't act >like a
disturbance.

I don't think you pondered this thought long enough. If a pre-calibrated act contributed to a higher-level control process, it could do so only with one setting of the higher-level reference signal. With a different reference signal, the precalibrated act would be too much, too little, or in the wrong direction. Remember that ACTS do not have consistent effects on OUTCOMES. The same act can have opposite effects. On one occasion of achieving a given outcome, a higher system might increase the output; on another occasion, it would produce the same outcome by decreasing the output. A chain mechanism can't do that.

I showed that there ARE reference levels in chain mechanisms: that point on the measurement scale of the stimulus at which no response is produced. But with or without reference signals, chain responses disturb things -- what difference does having a reference signal make?

>No, we should be looking for higher-level loops. And setting aside the
>claim that all control is continuous.

I'm still very unhappy with calling anything but negative feedback effects "control." Control implies a reliable repeatable outcome. In real environments, outcomes are subject to disturbance independently of the action of a behaving system. Only a negative feedback system can continue to produce the same outcome despite these disturbances.

By calling any other mode of action "control" you're throwing away the central phenomenon and going back to loose qualitative talk of the kind that has characterized the behavioral sciences for too long. I don't want to have anything to do with that.

If some organisms respond without feedback, then they must be very simple organisms living in very protected environments. Without this kind of limitation on the organism's abilities and this careful protection against disturbances, such systems would not be sufficient to allow survival to the age of reproduction. There might be such simple organisms and environments. I don't know. But if there are, those organisms do not survive by controlling anything. They survive in spite of not controlling anything.

If higher level control loops exist which modify lower-level open-loop responses, then control will exist at the higher level, but not at the lower level. This can work only if the inaccuracies and inappropriateness of the lower-level responses are unimportant. The vestibular-ocular reflex is very inaccurate and can't be adjusted very fast. The organism could get along without it altogether. It doesn't matter if the reflex causes the eye to miss the target by 20 degrees one way or the other - the oculomotor control systems will immediately make the eye direction exact. All the reflex does is make the locking-on slightly quicker.

When the outcome matters, control is continuous. You don't drive a car at high speed on a mountain road by looking out the windshield and giving the wheel a twitch once per second or so. If you did that your control bandwidth would be far too low and you'd go off the road or run into those Falling Rocks.

The biggest problem with allowing the behavior of open-loop systems to be called control isn't that there are no open-loop systems. It's that behaviors that actually couldn't happen without control, true feedback control, can be conceived of as open-loop behaviors and still be called controlling. The number of closed-loop phenomena erroneously interpreted as open-loop greatly exceeds the number of correctly identified open-loop phenomena.

The observer recognizes control by its effects: some outcome is produced over and over in a highly variable and uncooperative environment. So the observer correctly, intuitively, identifies the behavior as control behavior. But not realizing that feedback is required in order to explain the observed control behavior, the observer thinks that the same result could be achieved by an open-loop SR system. So the observer names the behavior in terms of its outcome: an "orienting response" or a "problem-solving response," and thinks an explanation has been found. In the observer's mind there's a notion that SOMEHOW there are stimuli that can produce the observed controlled outcome just by going through some magical black box. The observer doesn't see anything wrong with the idea that a stimulus input simply causes the outcome to occur. After all, the observer can see the connection. What do we need with all these complicated feedback loops anyway?

Look how easy it is to say that a threat causes the cockroach to run "away." Look how easy it is to be SATISFIED with that concept. If you can see the stomping foot and you can see the turn and the scuttling away, why isn't that enough to show that there's a simple response to a simple stimulus? Unfortunately it IS enough for many, many people. They aren't looking for closed loops. They have in mind some simple fuzzy connection diagram that goes between input and output, and because they haven't stopped to ask themselves in detail what those connections would actually have to accomplish, they don't see any difficulties. They assume that somehow the necessary circuits would exist and would accomplish what is observed.

Most of the names we have for behaviors are the names of outcomes, not actions. Even such a simple word as "walking" names an outcome. All that an organism with muscles can do is push, pull, twist, and squeeze. Everything else is outcome. When you realize this, you have to realize that there can't be any simple connection between an input and an outcome, even though it's easy to imagine simple connections between a stimulus and push, pull, twist, or squeeze. So when you imagine a direct connection between a stimulus and an outcome, you're implying far more complexity in the intervening processes than is obvious. If you're aware of that complexity, fine: if you're not, you're just waving your arms.

Oded Maler (920331) --

>Some of the arguments against S-R are not against the principle but
>against the practice of S-R psychologists.

True, but I have plenty against the principle, too: see above.

>The main argument in your post was a critique of "anthropormpization" >of perception and action, that is, characterization of events from the >point of view of an external observer rather than from the perspective >of the behaving organism itself.

I used the term "anthropocentrism," which is a little different. Anthropomorphism is OK when you're trying to understand a human being, although some might prefer "gynomorphism." It's inappropriate when you're trying to understand another species, because there's no guarantee that another species lives in a perceptual world like ours.

Anthropocentrism, on the other hand, means to me being centered in the observer instead of the behaving system, without realizing it.

>... if you climb up above the most basic sensation, even your models
>cannot escape from this problem as long as the categorization problem >is
unsolved. When you say that someone controls for being far away from >a
"dog", you can say *in principle* that "dog" can be somehow >represented by
internal perceptual coordinates of that someone, by >you'll never really
have such internal descriptions.

This is right, and it's always a problem. In principle we can apply The
Test to find out what perceptions another organism is controlling. But
finding candidates for The Test that aren't picked from the confines of
one's own perceptions is very difficult, maybe impossible without aid of
some sort. How would we recognize a controlled variable that we can't
perceive? The only answer I can see, and it's a very long-term one, is to
devise unbiased hypotheses in a random way, using a computer to generate
possibilities. I think the computer would have to be much smarter than the
ones we now have. It would have to be able to generate hypotheses about
perceptions of the type that the organism's nervous system would be capable
of creating through reorganization. Without such guidelines, I think a
truly random search for controllable perceptions would result in, as they
say, an NP-hard problem of the worst sort. I don't know how this problem is
going to be solved, or even if it can be solved. Perhaps we will always
have to be satisfied with approximations to true controlled variables,
particularly in other species, knowing that we can see them only as they
project into human perceptual space (and in particular, our own private
perceptual spaces).

>Even in this quote one can find a shadow of anthropormpization:

[I can't spell it either, half of the time]

>>This modified the sensor signal, and the response, which
>>had been aimed at one final state, is now aimed at a slightly >>different
final state.

In control theory, "aiming" isn't anthropomorphizing. It can be modeled in
terms of a control system with a reference signal that specifies the final
state. To say that a system's action is "aiming" at a changing final state
is only to say that the reference signal is changing during the action, and
control is changing the perceptual signal toward the changing reference
signal's value.

>If you just add time indexes (continuous, of course..) to r and s,
>and write
>
>r[t]=f(s[t],s*)
>
>then, by assuming that s[t] is somehow influenced by r[t'] for all >t'<t,
you can get an S-R formulation.

There is a subtle point here. What you say, strictly speaking, is true: the
current r is always AFTER the previous s. But it makes a difference how
previous the s is and how much r can change in that interval.

The problem with a strict formulation like r[t] == f(s[t-1]) is the
assumption that the response measure is instantaneous and can change by any
amount during one dt. If, however, the response measure can change only by

some amount delta during an interval dt, shrinking dt reduces the amount by which the response can change, proportionally. The limit is a continuous dependence of r on s at an INFINITESIMALLY SMALL time in the past, with the change in r tending to zero.

We're really talking about the fundamental theorem of the calculus here, which is the difference between a discrete and a continuous picture of nature. Zeno's Paradox arises from overlooking (or not yet having invented) this theorem and assuming that the tortoise can move from one position to another one half as far from the wall IN ZERO TIME. As soon as you think of the tortoise as moving with some finite velocity, the paradox disappears: the smaller the distance to be travelled, the faster it will be traversed. This converts the problem from a discrete, logical problem into an analog problem.

When responses change on a continuum (and any response involving physical movement has that property), we then have to ask how much difference there is between the current stimulus and the stimulus that "occured" dt ago. As dt shrinks, this difference approaches zero. But the response is still changing at a finite rate, so the change in the response from one dt to the next also approaches zero. In the limit, the response and the stimulus covary.

In real systems, of course, there is always a finite delay. But should we treat the behavior, then, as a series of closely-spaced stimuli and responses, or as an approximation to continuous change? Which way of representing the system, if either, gives the better representation?

In fact, treating the system as an approximation to a continuous one gives the better representation in almost all cases. The reason is the implication in the sequential analysis that the response is either present or not present, and in principle could go from present to not present or vice versa in ONE delay-time. In real systems, responses take many delay-times to build up after the stimulus appears, and many more to die out after the stimulus is gone. This aspect of behavior is missing from the S-R or discrete representation. The continuous representation, on the other hand, predicts exactly such gradual changes even if, in the finite-step approximation, the changes occur in steps.

Another test is to imagine dt going to zero and seeing what effect there is on the predicted behavior. In the discrete model, the prediction is that the changes in behavior must go faster and faster as dt is made smaller and smaller. In the continuous model, the changes in behavior approach a limiting speed which then stays constant as dt shrinks to zero. When dt is on the order of normal neural lags, the behavior predicted by a continuous model is already almost exactly like the real behavior; letting dt shrink to zero then makes no perceptible difference in the behavior of the model.

These considerations hold true in my model up to about level 6, control of relationships. At higher levels, however, discrete variables -- symbols -- come into being, and the kinds of behavior that occur are more like the discrete or sequential kind -- indeed, level 8 is defined as the sequence level. Now the S-R interpretation becomes more feasible and we have to look elsewhere than basic mathematics to discriminate SR predictions from closed-loop predictions.

>If the organism can tell the difference between changes in s caused >from

the outside and the same changes in s caused by its own actions, >it just means that you should replace s by a more refined perceptual >signal that can make these distinctions.

There isn't any way to make that distinction inside the system doing the controlling. A change in an input signal is the result of the sum of all influences on it: in that sum, the contributions of individual sources are completely lost. Exactly the same change could arise from the action alone, or from a smaller amount of action and an independent disturbance acting simultaneously. A higher system could sense the action (proprioceptively, tactily, or visually, for example) and compare the action with a copy of the lower-level controlled perception, and deduce the part of the perception likely to have been caused by independent disturbances. This would not help the lower-level system doing the controlling to tell which part of the controlled input was due to its own action. Neither could the higher-level system deduce how many independent disturbances were acting at once, assuming that their causes are not available to the senses (as they are usually not).

Fortunately, control systems do not have to know how much of a given perception is due to their own actions, or to sense the causes of whatever disturbances (in any numbers) are acting. The principle of control requires only sensing the state of the controlled variable itself, and producing actions that affect it. S-R systems, on the other hand are incapable of compensating for disturbances if they can't sense the cause of the disturbance. By definition, open-loop systems don't sense the effects of the disturbances on the outcome: if they did, they would be control systems and closed-loop analysis would have to be used.

Marcos Rodrigues (920331) --

>I think the foot-stamping situation is in closed loop. The cockroach
>action changes its perception of the environment, surely. My
>understanding of your theory is that an action modifying _physically_
>the environment is not required, provided that our perceptions of the
>environment change in direction of correcting the error.

You could be right, but I suggest that to make the cockroach's perception of foot-stamping into a controlled variable is probably too risky. You're quite right in reminding us that a controlled variable doesn't have to exist physically outside the organism (or to have a counterpart that does). Actually, I've been avoiding that point because the argument tends to drift off into epistemology, and it's easier to make the cases about S-R vs. control in terms of visible variables.

But yes, it's conceivable that the cockroach has a perception of a foot stamping, can recognize it as a significant (visual) event, and has an internal reference level set very low for such events. On the other hand, I think it's more plausible that the cockroach has only less advanced perceptions and has to control for simpler and more proximal variables. At least, if we can model behavior based on simple variables more directly connected with processes at the sensory interface, we'll be erring in the conservative direction, not atttributing higher levels to lower organisms until the data force us to.

>I can't see feedforward or open loop in the escape response.

Goody. Someone else on my side.

Best to all,   Bill P.

Date: Tue Mar 31, 1992 2:50 pm PST
Subject: categories, condition maintenance

Bill Powers (920331.1030)

> By calling any other mode of action "control" you're throwing away the
>central phenomenon and going back to loose qualitative talk of the kind
>that has characterized the behavioral sciences for too long. I don't want
>to have anything to do with that.

Which is why I tried to introduce the term `condition maintenance',
apparently without success.

Date: Tue Mar 31, 1992 3:01 pm PST
Subject: Re: Arm; S-R

[Martin Taylor 920331 17:00]
(Bill Powers et al. on the cockroach)

Very quickly--It occurs to me that Bill's argument that the cockroach
behaves so as to reduce the wind speed can hold only for the initial
turn and acceleration. After that, the approaching hand or foot will
be causing an ever increasing wind speed for which the cockroach cannot
compensate. The biggest wind blast will occur when the hand hits the floor,
just missing the scuttling cockroach. I don't see where the control system
can control in any way better than the ballistic control of a golf ball,
in this case. I do see it for the selection of the direction of escape,
and as Bill said, if we played golf for survival, not many would survive,
to which I add that those who did would be very fine golfers. Cockroaches
are very fine escapers.

Martin

Date: Tue Mar 31, 1992 8:56 pm PST
Subject: arm model V2

[From Joe Lubin (920331.1200)]

Bill Powers (920330.0800), Greg Williams

> Some solutions to the trajectory problem came together this morning.
> The model needs at least one more level, the transition level,
> which I now am convinced could also be called the "path" level.

I liked this idea when you first mentioned it as the "virtual target" in
your post of 920312.1000. It makes more explicit the notion of
continuous control (as opposed to relatively more ballistic components
of a motion).

In the post of 920312.1000 and in subsequent posts you speak of treating the alpha and gamma efferent reference signals totally independently. First of all there are two distinct types of gamma efferent: static and dynamic. Their different effects arise from the fact that they innervate different types of intrafusal fiber: static nuclear bag fibers and dynamic nuclear bag fibers with different mechanical responses to activation which gives rise to part of the static and all of the dynamic properties of the Ia afferent responses.

From what I have read, alpha and static gamma signals are usually proportional to one another: this is called alpha-gamma coactivation. In lower vertebrates, there is no separate gamma system and the spindles are efferently innervated by colaterals of the alphas. By separating the system into alpha (skeletomotor) and gamma (fusimotor) systems, more independent control can be exercised in adjusting sensitivity (adaptive gain control). I don't yet have a clear understanding of how the dynamic gammas are used, although it is clear that they are (i) silent for slow, mellow behaviors, and (ii) very active for fast or dangerous movements, and for imposed activity (say when a cat is picked up against its will).

Your switch to the alpha reference seems right to me, although separation of the two types of gamma signals might have taken care of your instabilities. Perhaps you could try to mask out the rate response (this is essentially what happens when static gamma is large and dynamic gamma is zero, although there always is some dynamic component to the Ia afferent signal) and try varying your alpha and gamma together and see if it works.

I think its a bad idea to conflate the extensor and flexor systems. Your reduction of the two contraction signals via (a2 - a1) does not capture all the salient aspects of the alpha commands (diagram of 920320.1100). This gives you no way of modulating stiffness/compliance. For a well-learned motion the agonist is contracted and the antagonist is inhibited rendering it nonresistent: this is called reciprocal innervation. For new learning of a motion or for a motion that requires stiffness in its execution, coactivation is used to tense both agonist and antagonist. In both strategies the same joint angle can be obtained, the difference lies in the stiffness of the joint. Reciprocal innervation allows for more rapid, fluid, less energy consumptive operation, and requires advanced knowledge of loads. Cocontraction is more exepnsive energetically but can deal better with unknown situations and heavier loads. I know that this particular modeling effort doesn't require the added complexity and I know that William of Occam would not be pleased with my suggestion, but I do think that most neuroscientists have difficulty abstracting from pure anatomies, so at least an explicit mapping from the anatomy to your circuit should be presented in your paper.

I just came across this:

> All this really requires modeling the
> opposing muscles separately.

> Maybe Joe Lubin and his students, plus Greg
> Williams, would like to carry this on to version 3.

Yup.

And from Bill Powers (920330.0800):

> I can see now that we will be able to make a
> prediction: the velocity profile, normalized as in the article, will
> also be the same for different spatial separations of initial and
> final positions, at least for rapid movements.

It's nice to be able to make predictions and have some data out there to test these. I'll send you a paper that references velocity profile invariance data for arm trajectories. Here is the abstract:

Neural Dynamics of Planned Arm Movements:
Emergent Invariants and Speed-Accuracy Properties During Trajectory Formation

Daniel Bullock and Stephen Grossberg
Psychological Review 1988 95:49-90.

Abstract

A real-time neural network model, called the Vector Integration to Endpoint, or VITE, Model, is developed and used to quantitatively simulate behavioral and neural data about planned and passive arm movements. Invariants of arm movements emerge through network interactions rather than through an explicitly precomputed trajectory. Motor planning occurs in the form of a Target Position Command, or TPC, which specifies where the arm intends to move, and an independently controlled GO command, which specifies the movement's overall speed. Automatic processes convert this information into an arm trajectory with invariant properties. These automatic processes include computation of a Present Position Command, or PPC, and a Difference Vector, or DV. The DV is the difference of the PPC and the TPC at any time. The PPC is gradually updated by integrating the DV through time. The GO signal multiplies the DV before it is integrated by the PPC. The PPC generates an outflow movement command to its target muscle groups. Opponent interaction's regulate the PPCs to agonist and antagonist muscle groups. This system generates synchronous movements across synergetic muscles by automatically compensating for the different total contractions that each muscle group must undergo. Quantitative simulations are provided of Woodsworth's Law, of the speed-accuracy trade-off known as Fitts' Law, of isotonic arm movement properties before and after deafferentation, of synchronous and compensatory "central error correction" properties of isometric contractions, of velocity amplification during target switching, of velocity profile invariance and asymmetry, of the changes in velocity profile asymmetry at higher movement speeds, of the automatic compensation for staggered onset times of synergistic muscles, of vector cell properties in precentral motor cortex, of the inverse relationship between movement duration and peak velocity, and of peak acceleration as a function of movement amplitude and duration. It is shown that TPC, PPC, and DV computations are needed to actively modulate, or gate, the learning of associative maps between TPCs of different modalities, such as between the eye-head

system and the hand-arm sytem. By using such an associative map, looking at an object can activate a TPC of the hand-arm system as Piaget noted. Then a VITE circuit can translate this TPC into an invariant movement trajectory. An auxiliary circuit, called the Passive Update of Position, or PUP, Model, is described for using inflow signals to update the PPC during passive arm movements due to external forces. Other uses of outflow and inflow signals are also noted, such as for adaptive linearization of a nonlinear muscle plant, and sequential read-out of TPCs during a serial plan, as in reaching and grasping. Comparisons are made with other models of motor control, such as the mass-spring and minimum-jerk models.

It appears that these people have thought about many of the issues with which you are grappling.

> Joe, How about some details on what your student has accomplished so
> far?

His first semester of work consisted only of getting you model to work, and creating a beautiful graphical interface on Silicon Graphics IRIS Workstations. He started with about 10 lines of your code and built it. He has just recently started his second semester of work. First he cleaned up the code so that one could implement a new level or a new ECS with ease. We are viewing this as a generalized vision/sensorimotor/control environment. I plan to use this for a few years. Our next steps are to implement (i) the lower levels (your muscle circuits) and (ii) make the visual depth computation less trigonometric. For this latter we are using a complex cepstral filter which operates locally on windows the size of ocular dominance columns in V1. This filter computes local retinal disparities.

Joseph Lubin

Date: Tue Mar 31, 1992 11:06 pm PST
Subject: arm model V2

[From Joe Lubin (920401.0100)]

Bill Powers --

In case my last post sounded like a bunch of complaining, let me make it clear that your circuit appears to be extremely insightful. I have not yet come across any formulation which provides such a clear justification for the anatomy. It helped me understand what might be going on.

I was rereading some of the Bullock and Grossberg paper. I think you'll really appreciate it.

Joseph Lubin