Date: Sat Aug 01, 1992 4:46 pm PST
Subject: sticks

To: Pat Alfano
From: David Goldstein
Subject: stick patterns
Date: 08/01/92

I don't have you post in front of me and am going on memory of it.

As I recall, you want to show that the stick patterns test is not measuring what it is claimed to measure, namely, something about spatial ability. You ended your post by asking: Am I being too ambitious?

What makes you say that it is not measuring what other people say it is? The different stategies which people use to do the task? This might be a good application for q methodology.

Write down each of the strategies on an index card. Select a group of people, perhaps researchers of this task, and ask each one to do a q-sort as follows: Rank order the strategies from those which most likely involve spatial ability to those most unlikely to involve spatial ability. Intercorrelate the q-sorts. Factor analyze the matrix of correlations. If more than one factor emerges, you have some support for your position. If only one factor emerges, you have some lack of support for your position. If more than one factor emerges, showing that people have different views on the matter, you could gain some sense of what each view is by looking at the qsorts for each grouping.

Date: Mon Aug 03, 1992 2:11 pm PST
Subject: PCT in Puerto Rico

[from Gary Cziko 920803.1700]

I thought it might be of general interest to CSGnetters to know that PCT will soon be making an appearance on the Commonwealth of Puerto Rico.

Joel Judd, the first person to do a doctoral dissertation based on PCT at the University of Illinois at Urbana-Champaign has got himself a job as a professor at the Interamerican University in San German. He will be involved in both teaching English and teaching about second language learning and teaching.

So you see, it IS now possible to both get a PhD related to PCT and find a job afterwards!--Gary

Date: Mon Aug 03, 1992 2:31 pm PST
Subject: Changing Addresses

[from Gary Cziko 920803.1715]

It has recently come to my attention that some people are on the CSGnet list with a different address from the one included on the return address of their posts (e.g., Bitnet address and Internet address, respectively). This can cause problems when an individual wants to change his or her status on the net.

I have therefore begun to change addresses on the list to make them consistent with return addresses. If I do this to you, you will get separate notices that you have been deleted from CSGnet (old address) and added to CSGnet (new address). You should not miss any posts in the process, but do let me know if you experience any problems as a result of the change.--Gary

Date: Tue Aug 04, 1992 2:03 am PST
Subject: astro report

Well, I've got a primitive version of astro going, tho I haven't done enough work with it to be entirely sure of what it does. Here's how it works (somewhat modified from my original posting on the subject).

Asto is a little spaceship who tries to track his mother ship using fore & aft thrusters (only 1-dimensional movement for now). Mother wanders back and forth across the screen, changing direction suddenly & at random (& when she hits the edges). Astro has a perceptual system that sets a reference level for velocity toward mother at:

    v_ref = c*x^(1/2)   (`^' = exponentiation)

where c is a constant settable by the reorganization system, whose optimal value is determined by the maximum acceleration that astro's thrusters can produce (when c is to big, there's overshoot, when it's to small, the tracking isn't aggressive enough).

Astro also has a intrinsic reference level for 0 distance towards mother, which drives `reorganization' (really `tuning' in a case like this), which is adding an increment (currently 0.05 or -0.05) to c. A running average of the intrinsic error is maintained, and every 100th iteration (100 chosen by trial and error, tho this value may well be too small), the current average intrinsic error is compared with that of the previous reorganization episode. If the new is greater than the old, the sign of the increment is reversed, otherwise the increment, is retained unchanged, and the increment is then added to c.

When c is too high, causing overshoot, the system seems to settle down, but when it starts near 0, it doesn't do so well, especially since once c gets to -1, then the intrinsic error climbs inexorably upward, but the increment just changes sign, and the problem never gets fixed.

This problem might be approached by having the increment vary with the magnitude of the average intrinsic error, but this would only reinforce the main intuition that I've gotten from this so far, which is that reorganization systems have to be rather carefully tuned to their domains in order to produce sensible results in a reasonable amount of time. (e.g. working critters need a lot of built-in knowledge about their environment, albeit perhaps rather high-level knowledge, such as about appropriate timescales for reorganization).

Another aspect of it that bothers me is that the reorganization episodes are sudden & discrete, while it would seem better to have it happening gradually, but determined by the long-term trend in intrinsic error. But I haven't figured out a plausible way to do this.

What next? Well, I suspect that astro is probably not worth developing much further, since the environment is too un-biological in nature (my original motivation was that since handling yourself with thrusters in a low friction environment is hard, a machine that could learn to do it would be impressive. But maybe it's too hard a problem, and too artificial).
So maybe I'll add coefficients of friction and turn him into a fish.

Avery Andrews@anu.edu.au

Date: Tue Aug 04, 1992 7:17 am PST
Subject: Re: PCT in Puerto Rico

That is the best news I have heard all summer. Good luck to Joel.

Date: Tue Aug 04, 1992 8:05 am PST
Subject: perception of rapid progress

[From: Bruce Nevin (Tue 920804 11:26:03)]

It seems a truism that people are not interested in alternatives (viz. PCT) so long as they perceive themselves as "getting somewhere" by their present means. Here's a perspective on a bunch of people who believe they're progressing toward their goals in robotics and artificial life. How could the news that they're not be framed in such a way that they see it as a contribution (mid-course correction) rather than a demand that they give up and do something else (abandon ship)?

If you receive cybsys-L mail, my apology for repeating the following:

From: Cliff Joslyn <cybsys@bingsuns.cc.binghamton.edu>
Subject: ALife III Conference Report, Hugo de Garis, ETL, Japan.
To: Multiple recipients of list CYBSYS-L
<CYBSYS-L%BINGVMB.BITNET@pucc.Princeton.EDU>

Really-Really-From: degaris@etl.go.jp (Hugo de Garis)
Really-From: alife@cognet.ucla.edu (Artificial Life Digest)

[The following is a cross-post from Artificial Life Research List Digest Number 081 Monday, August 3rd 1992. You are encouraged to subscribe by sending mail to alife-request@cognet.ucla.edu - Moderator]

Dear ALifers,
Here is a quick report on the highlights of the ALife III Conference which was held in Santa Fe, New Mexico in June 1992. It reflects the
personal biases and interests of the author.

Far and away the best paper was by Gerald Joyce, who spoke on his "evolution of
molecules" work. He takes RNA, and can mutate one, two or more specific bases, and then clone the mutants in huge numbers. These mutants are then subject to a selectionist test, so that only the more successful mutants survive.
The survivors are then further mutated etc, until molecules are evolved which perform some desired function. This is probable future Nobel work. Good luck Gerald.

John Koza used his Genetic Programming technique (evolution of Lisp programs) to evolve self reproducing systems, and told his audience that the size of the search space (with his primitives) was only of the order of a billion or so. This is exciting, because its a lot less than earlier estimates (eg von Neuman's
29 state reproductive cellular automata Turing machine). It means that it will probably be possible in a year or so to evolve selfreproductive systems which can also do something useful. This will be essential when nano tech is finally with us. Nanoscale machines will have to self reproduce in order to build macroscale systems.

The biological robot (biot) community went home with the message that an evolutionary approach to building biot nervous systems is the way to go. Brooks-style handcrafting has become too complex and needs an evolutionary approach. Even Brooks was talking about evolving gnat robots.

Dave Ackley showed that Lamarkian evolution can "blow the doors off" standard GA evolutionary learning. An amusing and engaging talk.

L. Buss and W. Fontana teamed up to present a mathematical theory of the development of systems towards life like behavior. They presented levels 0, 1, and 2, with increasing sophistication, e.g. self replication, then genomes, etc. It was hard to follow but felt important. At the end of the talk they claimed they had implemented it all in a computer program but they failed to present results (and this was to an audience of 70% computer types!).

A fast-talking half-crazy Canadian (M. Tilden) upstaged Rod Brooks by showing that he could do a lot of what Brooks does, but for a thousandth of the price. Tilden makes ultra cheap little robots with an amazing functionality. I hope Rod Brooks went home with the lesson. Funders take note!
Tilden will go down in ALife history for his reply to a non native english speaker's question, "Why do you talk so fast?". Tilden's reply was, "Whydoyoulistensoslowly?!"

J. Smits made tiny bilayer silicon strips curl up when a current is applied. He intends to use these strips as mechanical movers for "silicon ants".

Randy Beer now uses GAs to build his insect circuits. Darwinian Robots was one of the themes of the conference.

Maja Mataric (Brooks student) presented a video on Brooks version of swarm intelligence. 20 robots were supposed to perform behaviors by emergence. The video angles were too low to see effectively what was happening. Maybe she should have done the filming in a basketball court and taken the shots from a ladder.

Personally, I went away all fired up that "evolvable hardware" is possible.
I learned about FPGAs (software configurable hardware), so it will be possible to treat the configuration bitstrings of FGPAs as chromosomes in a GA. Thus the
technology may exist today to fulfill my dream of building "Darwin Machines".

I just got back to Japan, so this report probably reflects my jet lag.

Cheers, Hugo de Garis, Electro Technical Lab (ETL), Japan.

Bruce bn@bbn.com

Date: Tue Aug 04, 1992 8:26 am PST
Subject: Extremism or "Let us now ignore famous men"

[From Rick Marken (920804)]

Well, here it is; my first "post - meeting - post".

The Durango meeting was a joy, as usual. Every day was filled with helpful discussions and fascinating demonstrations.

A general impression from the meeting: many of us (people) have a great reluctance to abandon ideas that have at one time or another been deemed "important" or treated as "fundemental". Thus, for example, it is difficult to believe that "reinforcement" is a useless concept; that it is simply the result of a misperception of the phenomenon of control. A famous man said it was important -- and many famous men and women continue to use the concept -- so it MUST be important.

The same holds for other concepts that have become shibboliths of behavioral science because some now famous person once said it was important. Examples include: reflex, experimental method (IV-DV approach), statistical analysis, true score, intelligence, information and conditioning.

It is hard for people to believe that the fruits of the "great minds" of the past could be completely useless; many people think that there must be SOMETHING worth preserving from what went before -- even when they accept the value of a new point of view, such as control theory.

I find this respect for history admirable and I also think there actually might be some things worth preserving from the old point of view. They are the things that can move into the new control theory point of view effortlessly, with no one noticing. For example, the Copernican revolution held on to the Ptolmeic idea that there were planets moving in space. Similarly, control theory holds on to the idea that animal actions result from muscle contractions caused by efferent neural impulses. It just gets rid of the idea that these actions are caused (made to occur more frequently) by external events (reinforcers). Actions are, of course, part of a control loop; and some perceptual consequence of those actions is under control. The apparent effect of "reinforcers" is a result of looking at this control loop from the wrong perspective.

There are many places where control theory shows that some old, revered concepts are just plain wrong or irrelevant. Reinforcement a good example of a concept that is wrong; another example is the idea that properly executed "controlled" experiments (even if the results are perfect) reveal something about the psychological law that relates variations in the independent variable to those of the dependent variable. And another is the idea that behavior is "intrinsically variable" -- so that it it necessary to average over subjects or over trials with the same subject in order to estimate the "true" response.

Each one of these assumptions is believed important because it was pronounced by some "famous man" (yes, they were almost always men; another reason for exasperation); Skinner or R.A Fisher or whomever.

Thus, when a control theorist pronounces such ideas "wrong" there is some considerable consternation. Even when we make our case there is still the feeling that there must be SOMETHING that can be salvaged from the old ideas. But the fact is that control theory shows that these concepts are often just flat out wrong and can be safely (and profitably) ignored. This is the only proper response to these famous concepts. Unfortunately, the proper response is the one that makes the control theorist seem like an EXTREMIST (or, as one person put it at the meeting, a religious zealot).

I think the "extremism" label is just something one will have to put up with once he or she understands control theory. The alternative is to try to seem "moderate" by trying to find a place for "famous" concepts in control theory. Unfortunately, when this is done it usually results in compromises that, from a control theory view, are WRONG. For example, trying to study control using the IV-DV framework is wrong, not because a tenet of the PCT religion says "thou shalt not do experiments as the conventional psychologists do them", but because we know (and have shown in many experiments) that

o = - k.e (d)

What an organism does (its output, o) depends on environmental disturbances, d (not sensory inputs) -- and the relationship between disturbances and outputs depends on the nature of the environment that determines the relationship, k.e, between outputs and the sensory effects of disturbances. Thus, you are guaranteed not to learn about the nature of the controlling organism if you use the IV-DV approach. This is an "extreme" conclusion -- but the only correct one, I'm afraid.

There are certain control theory concepts that cannot be compromised with existing "famous" concepts -- without losing control theory. "Extremism" like this is, unfortunately, the price one pays for understanding control theory. I prefer to call it "integrity" rather than "extremism"-- but I'm sure it doesn't look like anything but "extremism" from the point of view of those who are not yet prepared to "ignore famous men".

Best regards Rick

Date: Tue Aug 04, 1992 9:54 am PST
Subject: Plooij references

[From Dick Robertson]
Here are the Plooij's refs. thzt I promised some of you at the conference:

1) Infantile Regressions: Disorganization and the Onset of Transition
Periods.
Plooij H C & F X Journal of Reproductive and Infant Psychology, Vol.
10,00-00,
1992. (2) Distinct Periods of Mother-Infant Conflict in Normal Development:
Sources of Progress and Germs of Pathology. Plooij, H C & F X. The Journal
of
Child Psychology & Psychiatry. In press (as of April, 92).

Best to all. Dick

Date: Tue Aug 04, 1992 11:40 am PST
Subject: Meeting report; think tanks; Astro

[From Bill Powers (920804.1100)]

The meeting is over and I've recovered. About 25 people attended, with Significant Others bringing the total to about 30. As usual, the flavor of the meeting was different from all the others. The general theme seemed to be the relationship of PCT to other fields. Wednesday night we did the schedule, taking about half an hour.

Thursday: Clark McPhail, Chuck Tucker, and Kent McClelland (our sociologists) spoke on teaching CT to their colleagues and students, referring to the Gatherings (Crowd) program and its reception, classroom approaches, and the merging of CT concepts into sociological topics such as power. Tom Bourbon and Dick Robertson spoke of their ventures in France and Belgium respectivly, teaching CT to students who were generally quite interested and attempting to teach it to various bigwigs who weren't much interested. Rick Marken gave an excellent demonstration of his three-level, six-system-per-level spreadsheet model. Greg Williams opened the subject of a possible PCT journal, to include Closed Loop plus refereed articles. As to refereeing, the idea was that if someone insisted on having a rejected paper published, it would be published along with the reviews (unless utterly atrocious).

Friday: Mark Lazarre (student of Tom Bourbon) described his master's thesis work on two-person learning of a cooperative control task. I did a show-and-tell of the new version of the arm model, a rubber-band version of the Cognitive Control model, and a reprise of the Bucket of Beans experiment for those who hadn't seen it before. Gary Cziko presented part of his compendium of Portable Demos, including one in which a large rubber-band was used to interfere with a subject bouncing a ball off a wall, showing that throwing does involve control despite its ballistic component. Dag Forssell went through some of his methods for presenting the basics of PCT to business executives. In the afternoon, Ed Ford showed a video tape of his PBS program. Friday night after the banquet, we had a 20-minute business meeting at which Wayne Hershberger was elected Vice-President of the CSG, to succeed to the presidency as usual next year (Chuck Tucker, elected to that post last year, automatically became President for 1992- 93). Mary continues as secretary-treasurer (what else?). Mark Olson presented his Master's thesis work on analysis of causation and control theory, ending up presiding over an hour-long discussion on angels, pinheads, language, and other related topics.

Saturday: Isaac Kurtzer (student of Tom Bourbon) presented his video analyses of a buck-of-rice experiment in which he obtained quantitative data showing the difference between the S-R version of the experiment (drop the rice in the bucket in one lump) and the continuous-control version (trickle the rice in over 10 seconds). Bill Williams presented his analysis of corporate interlocks (directors of major banks who sit on boards of major corporations where they can communicate), showing that in the world of Big Business, directors come face-to-face with directors of other corporations every few days, on the average (Hmm: potential explanation of some aspects of social control? He wouldn't say). His main point, however, was that everybody can learn to program if they have a good enough reason. Brian D'Agostino presented his doctoral thesis work on PCT and militarism, taking a lot of flak for his statistical approach but holding his own quite bravely and making some convincing points about PCT in this kind of research. Dick Robertson asked for advice in applying control theory to analyzing his own classes and their behavior as control systems controlling for grades. And finally, Wayne Hershberger asked what people meant by the term "Controlling for ...", setting off a discussion that went past closing time.

If I've forgotten someone, please fill in the blanks.

All these talks, as is our custom, were really more like discussions with one person standing up. Not a single person read a paper this time, although the students are given license to make more formal presentations if they want to. The non-hostile atmosphere is conducive to relaxed talks that are more like conversations.

Every evening after the evening session, most of the people retired to the lounge in the residence hall and resumed the discussion just for the fun of it; the last die-hards staggered off to bed at 2 or 3 AM, appearing somewhat bedraggled at breakfast at 7:30 AM but making it through the day in fine shape. Every year I wonder whether the group will maintain the kind of intensity it generated the last time; every year it does. Everyone talks to everyone else. This is the only interdisciplinary conference I know of where that works.

Bruce Nevin (920730) --

Your post on Studies in Artificial and Natural Intelligence was interesting, coming to me after the CSG meeting and contrasting so strongly with what goes on at that meeting. There are large numbers of people, it seems to me, who want to sit back and talk about understanding brains, intelligence, and so on, but who aren't willing to do much more -- they think they ARE doing something, I suppose. Armchair or think-tank science somehow impresses me less this week.

Avery Andrews (920803) --

Congratulations on venturing into the world of modeling. I think you've tackled a do-able problem. Let me offer a little help. It's a good idea to get a working control system before tackling reorganization. If you just set up your Astro problem a little differently, it will work fine.

The first thing to do is set up the physical situation. You have two objects in space, Astro and Mother Ship. Mother ship can move at some velocity in x, and so can Astro. Mother moves independently. Astro moves by using thrusters that exert a force in x.

Astro controls two relationship variables: velocity relative to Mother's velocity, and position relative to Mother's position.

You're postulating that Astro wants to perceive a velocity c*x^(1/2) relative to Mother. This poses a problem, because you can't take the square root of a negative number -- the program will blow up if x ever becomes less than zero. I suppose you're using the square root to cause slowing near the goal. But there's an easier way -- a control system will automatically slow down near the goal without any special precautions.

It doesn't really make sense to speak of a "perceptual system that sets a reference level for velocity..." because perceptions aren't reference levels. They're reports on the actual state of affairs, not the desired state. I think Astro will work better if you make the perception equal to the relative velocity: that is, vrel = vm - va, where vrel = relative velocity, vm = velocity of Mother, and va = velocity of Astro. You then have

```
                                 vrel*  (reference relative velocity)
                                   |       error, vrel* - vrel
                            ---->///// ------->-----
                  vrel---->|                        |
                        ///////                     |
                       +^   -^                      |
                        |    |                      v
                       vm    va <-----integral -- force
```

Now vrel = vm - va, with no square roots. This control system can control the velocity of Astro relative to Mother. It can't control position, however, because the velocity difference can be zeroed at any distance. This control system will be stable.

You need a higher level of control to control for position, using variations in the velocity reference signal vrel* as the means of control.

```
                  xrel*  (reference relative position)
                   |
             -    +v   error, xrel* - xrel
     ---------->//////-------->----
    ^                              |
    |                              k1    (gain adjustment)
    |                              |
    |                              |
    |                              | (reference relative velocity)
    |<--- xrel               -   +v       error, vrel* - vrel
    |                       ---- ///// ------->-----
    |             vrel---->|                        |
  \\\\\\                ///////                     k2 (gain adjustment)
  ^-  ^+                ^+   ^-                     |
  |   |                 |    |                      v
 xa   xm <--integral-- vm    va <-----integral -- force     --
 ^           MOTHER          |         ASTRO                  |
 |                           |                                | (Environ
  ----<---integral-----<-----                                 |   ment)
          ASTRO                                             --
```

The variables xm and xa represent Mothers position and Astro's position. xrel is the relative position, xm - xa. xrel* is the relative position that Astro wants.

Mother's velocity is arbitrarily set (vm). Mother's position, as a program step, is xm = xm + vm*dt, where dt is set to some small value like 0.01. This integrates velocity to position. You can make Mother bounce off the ends of the screen by reversing vm.

Astro's velocity is the integral of the applied force divided by Astro's mass:

va = va + (force/M)*dt

Astro's position is the integral of Astro's velocity:

xa = xa + va * dt.

Thruster force is proportional to Astro's velocity error:

force = k2*(vrel* - vrel)

And of course vrel is just the same as va unless you want to play with perceptual functions of other kinds.

For position, we have simply

xrel = xm - xa

where xrel = relative position
      xm = Mother's position
      xa = Astro's position

The velocity reference signal vrel* is proportional to the higher-order error signal, so

vrel* = k1*(xrel* - xrel).

If you want Astro to seek the same position as Mother, just set xrel* to zero. If it's nonzero, Astro will tag along at a fixed distance.

In each part of the diagram, there's a potential scaling factor. You could make one or more scaling factors the target of reorganization. The intrinsic variable that could be monitored might be total error at both levels of the diagram (with an intrinsic reference level of zero).

If you use the constant k1 as the target of reorganization, you don't want to vary k1 itself at random. That will give large jumps in performance and there's no guarantee that you won't blow the system up.

It's better to use a delta that goes with each target of reorganization. Delta is a small number between, say, -0.1 and 0.1, chosen at random. This number is added to k1, the target of reorganization, on every iteration, so k1 is continually changing. But wait ...

The intrinsic error doesn't depend on the sign of the difference between the intrinsic variable and its reference level, but only on the magnitude. So you can use |i*-i| as the error (in this case, |-i| because the intrinsic reference level i* is zero). Then you compute the rate of change of i, which just i(now) - i(previous). If this rate of change ever becomes greater than zero, you institute a reorganization, which amounts to picking a new value of delta at random, in the range -0.1 to 0.1. Between reorganizations you just keep adding this delta to k1, the constant that is being reorganized.

You may want the changes to get smaller as k1 approaches the optimum value (assuming there is one). To do this, you simply multiply delta by the absolute intrinsic error before adding it to k1. As the error gets smaller, delta will decrease in size, going to zero when intrinsic error is zero. You're still reorganizing, but the amount of reorganization of k1 becomes zero.

So: if you choose total error as the intrinsic variable i,

```
i = |vrel* - vrel| + |xrel* - xrel|

if (i(now) - i(last) < 0)  delta = 0.2*(random - 0.5);    (reorganize)

k1 = k1 + 0.001* i*delta;     (change k1 on every iteration)
```

The random function is assumed to return a value between 0 and 1. I put the 0.001 in the last line to indicate that you'll want to make the changes rather slow -- how slow is a matter for experiment.

You can set up control systems for Astro in one, two, or three dimensions. Each one will be a one-dimensional two-level system as above, operating in x, y, or z. You need only one reorganizing system, however -- just have it sense the sum of the six absolute error magnitudes (two in each system). When a reorganization is called for (rate of change of intrinsic error is positive), just pick new values for three deltas, one for each axis. Use the random function for each one individually, so the deltas will vary independently.

I'll leave it up to you to get this to work. You shouldn't have any trouble with the control systems. The reorganizing part might get tricky.

Best to all, Bill P.

Date: Tue Aug 04, 1992 8:27 pm PST
Subject: astro

[Avery Andrews (920805)]
Bill Powers, <timestamp losgt> --

Most of what you suggest is already present and working. a run-thru:

Diagram one: velocity control system, there and working (I thought it was kind of neat that you can use thrusters to control velocity without knowing anything about second derivatives).

Diagram two: astro currently has an `implicit control system' for proximity to mother, with reference level 0 (whence my remark about the perceptual system setting a reference level for velocity: if xrel* is fixed at 0, the xrel perceptual signal determines the reference for velocity. But I think I will set up an explicit position-control system, and explicity drive reorganization with its average error (this what is happening implicitly at the moment).

What then of the gain adjustment k1? Here things get a bit tricky. The first version of astro used a linear gain function, so that vrel* = -xk. But this works rather badly, because the accelerations it calls for are too high when x is large, too small when x is small (if you move at the controlled-for velocity, the resulting accelerations will be a linear function of distance from mother). Reviving a bit of my (seriously rusty & dormant) calculus seemed to reveal that making the reference velocity proportional to the square root of the distance would give constant accelerations, which is what you want with thrusters.

The actual gain function involves a bit of sign-sillybuggery to deal with square-roots of negatives; technically it's:

```
  f(x) = 0  if x=0
  f(x) = c*x^(1/2) if x > 0
  f(x) = -f(-x) if x < 0
```

A bit messy and hoc, but maybe it would look better if I set up distinct control systems for the fore and aft thrusters.

Reorganization:
randomly choosing a new delta when the average error increases is possible, but maybe not very effective in this one-dimensional system: what really matters about delta is its size-range and its sign. If things are getting worse, you don't want another delta in the same direction, and there is only one other direction to go on (unlike the e-coli case, where there are an infinite number of different directions to go in, with nice theorems relating the direction of travel and the consequent rate-of-change-of-concentration of nice and nasty substances). So I think that sign-flipping is probably better than random choice. Scaling the size of delta to the average error looks like a good idea too, tho the whole thing runs so slowly that it takes a while to figure out if it works.

With fixed delta & sign-flipping, c seems to slowly drift downward to a sensible rate, but at an agonizingly slow rate; with delta proportional to absolute error, things seem a bit quicker. But I think a really need a statistical method for evaluating the thing, given it's current level of performance ...

Avery.Andrews@anu.edu.au

Date: Wed Aug 05, 1992 5:29 am PST
Subject: re: PCT in Puerto Rico

Great News!! Congratulations and bona fortuna!

Bill C.

Date: Wed Aug 05, 1992 6:32 am PST
Subject: Re: perception of rapid progress

From: Eric Harnden (920805)
Re : Bruce Nevin (920804)

>It seems a truism that people are not interested in alternatives (viz.
>PCT) so long as they perceive themselves as "getting somewhere" by their
>present means. Here's a perspective on a bunch of people who believe
>they're progressing toward their goals in robotics and artificial life.
>How could the news that they're not be framed in such a way that they
>see it as a contribution (mid-course correction) rather than a demand
>that they give up and do something else (abandon ship)?

my immediate reaction to this, i'm afraid, was to rant a little. i will not, both for the sake of propriety, and to defend myself at the outset from what unfortunately i imagine would be the response... something along the lines of 'see... resistance to new ideas is so strong that a suggestion of change brings forth only anger, not re-examination.'

first... what does mr nevin suppose the goals of artifical life research to be, exactly?

second... i think that there is a significant difference in the usage and examination of 'behavior' from the points of PCT and artificial life. i'm not sure how to express it, but it occurs to me as a difference between thought and action (sorry for the trite terminology, but i'm not feeling particularly articulate right now). i found mr powers' statement in BCP that he was after, not a model of behavior, but a model of how an organism might be organized that it may behave, as a fundamental. both an insight and a demarcation as to what PCT's realm of competence is. Kuhn, you may remember, says that one of the functions of a science's self-declaration as a science is to *limit* its universe of discourse... to provide guidelines for what it is *not* about. as an interested observer and dilletante, i see no particular conflict between CT and Alife that should require one to impose on or seek out the other at this time.

third... occam is frequently invoked here. CT claims to offer better explanatory power than the current mainstream of psychology, with less recourse to mumbo. 'does it work' is the mantra, as it should be in a field that derives, as CT does, from an engineering perspective. well then, the razor cuts both ways. Alife research is developing models that function. tools are being made which accomplish computational tasks. i really don't see any reason why its practitioners should modify their approach, since they are in fact (not just appearance) progressing toward their goals.

-----------< Cognitive Dissonance is a 20th Century Art Form >-----------
Eric Harnden (Ronin)
<HARNDEN@AUVM.BITNET> or <HARNDEN@AMERICAN.EDU>
The American University Physics Dept.
4400 Mass. Ave. NW, Washington, DC, 20016-8058

(202) 885-2748
---------------------< Join the Cognitive Dissidents >-------------------

Date: Wed Aug 05, 1992 9:27 am PST
Subject: Congrats; Astro; Alife

[From Bill Powers (920805.1000)]

I add my congratulations to Joel Judd (and his wife, who should enjoy a country where she can use her FIRST language). A new center from which PCT can grow, and a competent gardener to manage the process.

Avery Andrews (920805)--

Great! I don't object to using position as the intrinsic variable for reorganization; it's as good a way of exploring reorganization as any.

>What then of the gain adjustment k1? Here things get a bit tricky. The first >version of astro used a linear gain function, so that vrel* = -xk. But this >works rather badly, because the accelerations it calls for are too high when x >is large, too small when x is small (if you move at the controlled-for velocity, >the resulting accelerations will be a linear function of distance from mother).

Something's wrong here, or miscommunicated. If vrel* = xrel* - xrel, it's true that velocity will be a linear function of position error, but acceleration should vary in a way that brakes the motion to a stop just as xrel becomes equal to xrel*. If you start from a standstill, the initial acceleration should be toward Mother, and then reverse as the distance begins to decrease. If all this happens too slowly, you probably need to adjust k1 and k2 (in my second diagram). I'm confused when you talk about "moving at the controlled-for velocity." The controlled-for velocity should be large when the position error is large, and decrease as the position error decreases. Increasing the velocity reference signal should result in a positive acceleration; decreasing it should result in a negative acceleration. The velocity should change right along with the velocity reference signal.

Reorganization:

>If things are getting worse, you don't want another delta in the same
>direction, and there is only one other direction to go ...

That's fine for the one-dimensional case, but then reorganization isn't blind and you can't extrapolate it to the general case. It's OK if a random choice moves you in the wrong direction, because it will be followed immediately by another reorganization -- the error will get larger. There's only a 1 in 16 chance of making 4 wrong moves in a row, and the parameter won't change much because the steps are small. Once the direction is right, you'll get many iterations before another reorganization is called for, and will quickly make up the lost ground.

Want me to try this to compare with your approach?

I don't quite picture where you're putting this square-root thing into the control diagram.

Eric Harnden (920805) --

I'll let Bruce Nevin reply for himself, but I'd like to ask you a couple of questions.

What do Alife people consider a living system to be? That is, is there any basic principle of behavioral organization that underlies their attempts to explain and imitate the behavior of living systems?

Do Alife people assume that once a cognitive command to act has been emitted, the physical actions that will realize it simply follow from the command? You've been around CSGnet long enough to know that we see behavior -- that is actions -- as a variable means of accomplishing repeatable outcomes in a disturbance-prone environment; that it is perception of the outcome that is under control, not the action that helps to bring it about. Do Alife theories that relate thought to action take this closed-loop phenomenon into account?

Do Alife people see the process of computing behavior as assessing the environment, computing the actions that will achieve a given result in that environment, and then executing the actions? This is just another way of putting the preceding question, but it's more specific to some of the behavioral models that seem to come out of Alife.

I agree with you that we PCTers are not too careful about making inflammatory remarks. You have a right to bitch about them, even if only from the standpoint of demanding equal respect. If everyone will be reasonable and respect the will of others, we can examine differences in our approaches and work them out. I don't mean to rule out sharp remarks and retorts -- it's just that we'll get further faster if we keep in mind that being friends is better than being enemies. You can say pretty much anything you please to me as long as I can be sure you don't hate me. I hope you feel the same way.

As you say, Alife researchers have goals and are heading toward them. This all comes down to what's meant by designing artificial life, as opposed to just designing a system that will accomplish something. Maybe you could say something about that.

Best to all,   Bill P.

Date: Wed Aug 05, 1992 9:45 am PST
Subject: Re: perception of progress

[From: Bruce Nevin (Wed 92085 12:21:44)]

(Eric Harnden (920805)) --

I apologize if I have offended you. I imagined (wrongly) that there were no ALife insiders on the list, and consequently framed my hasty note exclusively in terms of a PCT perspective. I would have put more effort into bridging had I been more thoughtful.

It seems that we both understand that ALife folks in the cross-posted message believe they're progressing toward their goals in robotics and artificial life. It sounds like you take strong exception to the proposition that they in fact are not.

It may be that I do not understand their goals, as you suggest. A succinct statement of those goals from you would surely not hurt, and I am always willing to learn.

My understanding is that their goals in a general way include producing computational or robotic entities that act like living things.

What I see is attention to behavioral outputs of entities and populations of entities from the observer's point of view. While there are problems with the status of social conventions, a central PCT insight is that behavioral outputs are incidental byproducts, variable means for attaining consistent ends in a disturbance-prone environment. Since I see no sign of that insight in the ALife literature that I have looked at (and I am admittedly very much an outside observer), I expect that they will discover that despite rapid initial progress with concepts and computational tools, their efforts will lead to products that are in fundamental ways unlike living things. (The obvious parallel is to early conceptual progress in AI followed by very limited implementations that ultimately make good only a fraction of what was promised. I don't put much weight on this parallel, only to note that it is there.)

I believe ALife efforts will lead to products that are in fundamental ways unlike living things because it appears to me extremely likely that living things are perceptual control systems, as PCT claims, yet I see no indication that ALife entities control input perceptions relative to internally-held reference perceptions by way of effectors through the entity's environment, not even as emergent properties of those entities. Hence, the claim that they are not really achieving the goals that they think they are (entities that are like living things, or even that act like living things consistently amid disturbances).

If you can show that I am wrong in the second part of this (the claim about ALife), then rapprochement should be easy. If you cannot show that this is wrong, then I see two choices for you (there may be more). You might seek to reconcile the conflict by deciding that you must persuade us that the first part of this is wrong (the claim about living things being control systems). Or you might decide that persuading a bunch of fanatics, who are on the fringes of where the real action is anyway, is not worth the trouble, and just go off and forget about PCT.

Hence my question:

>How could the news that they're not [really achieving their goals]
>be framed in such a way that they
>see it as a contribution (mid-course correction) rather than a demand
>that they give up and do something else (abandon ship)?

You were groping for a way to articulate

>a significant difference in the usage and
>examination of 'behavior' from the points of PCT and artificial life.

You suggested (with apology) "thought" vs "action," and then suggested a parallel with Bill Powers' distinction between a model of an organism's behavior (behavioral outputs) and a model that behaves as the organism does. Then you take this as "a demarcation as to what PCT's realm of competence is," with reference to Kuhn's discussion of the delimitation of a science (I used the metaphor of the magic circle for this in recent discussion with Bill).

I see only one way that you can conclude from all of this that there is "no particular conflict between CT and Alife that should require one to impose on or seek out the other at this time," and that is that you must understand ALife as modelling behavioral outputs; or at least you must believe that ALife does not aim to create models whose behavior (as byproducts of control) is like that of the organisms modelled. But this is just the claim about ALife that I invited you to demonstrate was wrong.

Perhaps I have not understood you at all. Please explain what this means:

>Alife research is developing models that function.

Thanks for your help with this.

Bruce bn@bbn.com

Date: Wed Aug 05, 1992 11:17 am PST
Subject: Re: perception of rapid progress

(From Rick Marken (920805)

I would like to add my voice to those of Bill Powers (920805.1000) and Bruce Nevin (Wed 92085 12:21:44) who asked Eric Harnden (920805) to explain what ALife is about. Bruce wrote a particularly articulate reply to Eric, who seemed to take umbrage at Bruce's implication that, given its current assumptions, ALife was not going reach the goal of producing models that act like living systems.

Bruce said (to Eric):

>I believe ALife efforts will lead to products that are in fundamental
>ways unlike living things because it appears to me extremely likely that
>living things are perceptual control systems, as PCT claims, yet I see
>no indication that ALife entities control input perceptions relative to
>internally-held reference perceptions by way of effectors through the
>entity's environment, not even as emergent properties of those entities.
>Hence, the claim that they are not really achieving the goals that they
>think they are (entities that are like living things, or even that act

>like living things consistently amid disturbances).

>If you can show that I am wrong in the second part of this (the claim
>about ALife), then rapprochement should be easy. If you cannot show
>that this is wrong, then I see two choices for you (there may be more).
>You might seek to reconcile the conflict by deciding that you must
>persuade us that the first part of this is wrong (the claim about living
>things being control systems). Or you might decide that persuading a
>bunch of fanatics, who are on the fringes of where the real action is
>anyway, is not worth the trouble, and just go off and forget about PCT.

I agree with Bruce's analysis and would very much like to see a reply to it by someone who is really involved in ALife, as Eric seems to be.

Bruce again:

>I see only one way that you can conclude from all of this that there is
>"no particular conflict between CT and Alife that should require one to
>impose on or seek out the other at this time," and that is that you must
>understand ALife as modelling behavioral outputs; or at least you must
>believe that ALife does not aim to create models whose behavior (as
>byproducts of control) is like that of the organisms modelled. But this
>is just the claim about ALife that I invited you to demonstrate was
>wrong.

Beautifully said!

If Eric would answer Bruce's (and Bill's) questions about ALife it would, indeed, help all of us who are trying to understand how living systems work (I assume that that includes most of us on csg-l).

Thanks Rick

Date: Wed Aug 05, 1992 5:48 pm PST
Subject: astro, Alife

[Avery Andrews (9208060] (Bill Powers 920805)

Square roots:

The linear position -> reference relationships works, but doesn't make efficient use of the thrusters. Intuitively, the reason is this: when you are far from the goal, the reference velocity is high. So you if you are moving at it, it is also changing fast (since the reference velocity is a linear function of position, and you are moving fast). But when you are near the goal, you are moving slowly, so the reference velocity is changing slowly. So although the target velocity is a linear function of position, it is a nonlinear function of time, with a steep slope far from the goal and a shallow one nearby. So the acceleration, which is what the thrusters produce, will be way beyond their maximum capacity far away, and way below it nearby (and the results look pretty dumb on the screen). I in fact have a calculus argument the acceleration will be k1^2*x, where k1 is the gain factor in the second diagram, but there's a step I'm unsure of in the argument).

Motivating the square-root taking: suppose you want to get from point A to point B in minimum time (assuming newtonian physics), and energy expenditure & avoiding obstacles is not an issue. Then you want to go at maximum acceleration on the first half of your trip, and maximum deceleration on the second (the time-position plot of the second. half will look like that of a falling coffee-cup, with time reversed). So, for x = position, v = velocity, and t = time, we have:

```
x = k*t^2    (i)
x = v*t      (ii)
```

where k is determined by the acceleration, and we have cleverly chosen our coordinate system so that x and t are both 0 at the moment of arrival.

So, solving for v as a function of x, we get:

```
t = (x/k)^(1/2)

v = t/x

v = k1*x^(1/2)
```

Of course, it would be interesting to get the exponent to be selected by a reorganization process, but since it follows from fixed laws of nature, it would be sensible for it to be hardwired (but maybe I will experiment with getting it set by reorganization as well). As for where it goes, I'm currently putting it into the function that turns a position error signal into a velocity reference level, but maybe it could be put into the relative position perception system as well.

Deltas:
I agree that when reorganization is multidimensional (one intrinsic error controlling change in several control system parameters), something like the method in Bill's simultaneous equational-solver is needed - I just want to get a feel for the 1-dimensional case first. The point is that what you want is random choice of direction, rather than of the total magnitude of delta.

Alife:
I'm not at all happy with the way the Alife discussion is going. Diehard PCT-ers believe that living systems have to be made out of control systems, but nobody really *knows* this yet. The only way to actually find out is for models based on this idea to be compared with others that aren't, and this is exactly the opportunity that Alife provides. But PCT-ers will have to show up at the party bearing goodies, rather than stand outside sneering at it.

What seems to me to be the most compelling criticism of Alife is that the success of their models is judged by their producing behavior that `merely' looks lifelike, rather than fitting any principled specifications of what lifelikeness is. But our judgements of the behavior of lifelike agents are not just an exiguous construction based on our current culture, but the end result of 600 million years or so of competitive R&D by multicellular organisms trying to fake each other out, and so are not to be taken lightly, in my opinion.

A line I am thinking along goes like this: what the nervous system of a living creature is is a continuous vector transducer whose outputs influence its inputs. The expected behavior of such a system is a non-viable combination of chaos and catatonia (outputs being forced to extreme values by positive feedback effects), and the (putatively) only way to avoid this consequence is to load the systems to the gills with negative feedback loops. So the PCT perspective ought to improve the design of Alife systems, and the analysis of Rlife ones (organisms).
I think this ought to be an adequate rationale for the pursuing the PCT approach within Alife circles, but it's got to be followed up by flash demos. Gatherings (nee Crowd) looks like a fine start in this direction, and one of my motivations for astro was to see if one could construct a system that could learn to handle itself with thrusters in a frictionless environment, which is a rather challenging problem (as I discovered while playing with one of my kids' toy hovercraft).

Avery.Andrews@anu.edu.au

Date: Wed Aug 05, 1992 8:35 pm PST
Subject: Die Hard III

(From Rick Marken (920805.2000))

Posting from home so here is some leader...

Avery Andrews (9208060) says:

>I'm not at all happy with the way the Alife discussion is going.

Actually, it seems to me that it hasn't started yet.

>Diehard PCT-ers believe that living systems have to be made out
>of control systems, but nobody really *knows* this yet.

This particular die hard PCTer does NOT believe that living systems "have to be made out of control systems". Never have believed it; never will. What this PCTer does "believe" is that living systems CONTROL -- it's a PHENOMENON. Try the test; you apparently have some living systems running around your house. Push on some variables -- watch those little systems control. Right now, the only theory that we know of that can control is CONTROL THEORY. If another theory comes along that can account for the details of control as exhibited by living systems better than control theory, then I'll believe that one. I'm fanatic -- but I'm flexible.

>The only
>way to actually find out is for models based on this idea to be
>compared with others that aren't, and this is exactly the opportunity

>that Alife provides.

This is not quite true. We have shown over and over again that cause-effect oranizations (the basis of all current non-control theory models of behavior that I know of) cannot control. The problem is that non-control models can often produce non-control phenomena rather nicely. So when you compare the models (in a situation where it is impossible to see that there is no control) it seems like a wash -- or the more detailed non-control models appear more impressive. The only "play-off" between models that makes sense (from a PCT point of view) is one where behavioral outcomes are produced in the face of unpredictable (and undetectable) disturbances. I've done such a playoff with psychological models (as has Tom Bourbon and Bill Powers), to the total lack of interest of the community of psychological theorists. Sometimes winning just ain't enough.

>But PCT-ers will have to show up at the party
>bearing goodies, rather than stand outside sneering at it.

Any sneering comes from the exasperating tendency of "non-PCTers" to ignore the most important "goodie" that die hard PCTers have to offer -- the phenomenon of control. Since students of living systems seem determined to remain uninterested in trying to figure out what the hell they are trying to explain (since they already "know" what's to be explained -- BEHAVIOR, of course) then it does seem that we will have to try to point out the beauty of control theory with other "goodies". Bill Powers' "Little Man" demo, for example, shows how computationally efficient it can be to build a behaving system (forget that it also controls) using the principle of control of perception. By doing so, you eliminate the need for gargantuan computations of inverse kinematics and ridiculous assumptions about what the system can know about the state of the environment in which it is producing results. The roboticists will love this "goodie" (I bet -- based on Tom Bourbon's report of the reaction to it at Aix) but it will still take some time to show that control of perception is not just a good technical strategy; it is also the strategy that is used by living control systems because there is no other way to produce consistent results in a disturbance-prone environment (ie. to control).

>What seems to me to be the most compelling criticism of Alife is that
>the success of their models is judged by their producing behavior that
>`merely' looks lifelike

That's the ONLY criticism; another way of saying this is " their models don't control." (probably -- but I'd be happy to be proved wrong. Indeed, that's what Bruce and Bill and I want to learn from the ALife dicussion; are Alife models trying to explain control? But you say you don't like the direction of that discussion. I can't think of a better one).

> But our judgements of the
>behavior of lifelike agents are not just an exiguous construction
>based on our current culture, but the end result of 600 million years
>or so of competitive R&D by multicellular organisms trying to fake
>each other out, and so are not to be taken lightly, in my opinion.

Darn. I always have to look up the word exiguous.

>A line I am thinking along goes like this: what the nervous system of
>a living creature is is a continuous vector transducer whose outputs
>influence its inputs. The expected behavior of such a system is a
>non-viable combination of chaos and catatonia (outputs being forced to
>extreme values by positive feedback effects), and the (putatively) only
>way to avoid this consequence is to load the systems to the gills with
>negative feedback loops. So the PCT perspective ought to improve the
>design of Alife systems, and the analysis of Rlife ones (organisms).
>I think this ought to be an adequate rationale for the pursuing the PCT
>approach within Alife circles

I think you are putting the mechanisms before the phenomena. You assume that living systems exhibit all kinds of behaviors that need certain mechanisms to fix them up. I don't believe there is any evidence that living systems normally exhibit these behaviors -- at least not nearly as much evidence as there is that organisms control.

But why work so hard at trying to find a place for PCT in Alife? Let's see what Alife is about and maybe it can get along just fine without us fanatics (as Bruce Nevin correctly suggested).

Best regards   Rick (Bruce Willis) Marken

Date: Wed Aug 05, 1992 10:03 pm PST
Subject: Die Hard 3

[Avery Andrews 920806.1549]   (Rick Marken (920805.2000))

>I think you are putting the mechanisms before the phenomena. You
>assume that living systems exhibit all kinds of behaviors that need
>certain mechanisms to fix them up. I don't believe there is any evidence
>that living systems normally exhibit these behaviors -- at least not nearly
>as much evidence as there is that organisms control.

I think you're missing my point here. Which is that given what a nervous system in an environment *is*, its expected behavior is to exhibit the symptons of grand mal epilepsy. The fact that this condition is actually rather rare demands an explanation, which PCT seems to offer.

Avery.Andrews@anu.edu.au

Date: Thu Aug 06, 1992 3:35 am PST
Subject: Alife and all that

[Oded Maler 920806]

I don't want to get into all that, just correct an apparent misconception that some CSG-ers seem to have: behavior-based robotics (a-la Brooks) is just one branch (not the main one) in the Alife coalition. The mainstream is more related to things

such as cellular automata, self-replication, emergent proerties,
self-organization, complex dynamics and this kind of stuff.
I also think that the "crowd" demo is the PCT work which is closer
in the spirit to some Alife research.

Of course AL shares *some* of the stupidities of AI (but refrains from
others [and has its own special ones {but *every community* has its own,
isn't it?}]).

And in this spirit, since nobody objects to casual forwarding,
I enclose something I found today incidently in sci.chem,
concerning simulation of living cells.

--Oded

"I was interested to read that someone wants to simulate the chemical reactions occuring inside a cell. As this is my field of interest I thought I might make a few comments. As a number of people have already (and correctly in my view) suggested, the short answer to this is no you can't easily simulate a cell. Forget about the need for a big computer, that's a minor problem compared to the lack of data, i.e real data on the kinetics of the enzymes that make up your cellular model. As someone said in their reply, people have been trying to simulate pathways for *years*, even before the digital computer was invented, pathway simulations were being done on analogue computers. The late David Garfinkel was a pioneer in this field. His largest model must have comprised of up to 400 reactions. Of course lots of other people (including myself) have since tried to simulate bits of cells. I pretty much abandoned the whole idea though for two reasons. One is that there is so little relevant data on enzymes that the model one ends up with bearly resembles the real thing, and secondly if you do end up with a model that manages to simulate a cell or some part of a cell, you have the problem of interpretation. These systems can get so complicated (it's not like your simple pendulum or inclined plane as in physics) that one doesn't have a clue for example why a particular metabolite should be going up in concentration while another is going down, and I'm just thinking here of simple steady state behaviour let alone transients, periodic or chaotic behaviours. There is also the comment that if the model is as complicated as the real thing then one might aswell study the real thing (unless of course you just like simulating things with a computer).

I suppose it really depends on what one wants to get out of the model. If you are building a model which you hope will reproduce how a real cell behaves that I think you're in trouble (for the already mentioned reasons), of course go ahead and have a go but you'll probably see what I mean after a while. The other reason why you might want to simulate pathways is to try and understand basic pathway behaviour. In other words what sorts of behaviour can pathways elicit. At the most basic level pathways can be in one of three states, either equilibrium (not very interesting), steady state (including unstable steady states) or transients, i.e moving from one state to another. So first of all one has to decide what state one is going to investigate. Secondly, something others have mentioned, are you going to assume a homogeneous system (i.e no concentration gradients and therefore no PDEs, just ODEs) or a heterogeneous system. Clearly a full PDE model (even in only two dimensions) requires a big fast computer. This still leaves the problem of interpretation of the results. Until a few years ago there was little or no formal theory which could describe metabolic pathways. Of course there is enzyme kinetics (a formal theory of enzyme behaviour), but trying to extend this to whole pathways (even to a two step pathway) is a hopeless task. The equations are simply insoluble (ok a two step pathway is soluble but the solution is about 3 pages of A4 long!). So where does one go from here? Since I abondoned realistic simulations I have turned to Metabolic Control Analysis (MCA) (or a similar theory called Biochemical Systems theory). These may sound a bit wooly, but believe me they are not. They have probably made more of a contribution to our understanding of pathways than all the simulations that have gone before.

The essential goal of MCA is to be able to relate the genotype of an organism to its phenotype. This is of course the classical objective in genetics (I really mean genetics here not cloning) and two of the originators (Kacser and Burns) were geneticists. The two other originators (Heinrich and Rapoport) were biochemists. It would be difficult to properly discuss MCA here (if there is any interest I'll mail a proper summary to the net) as this mail is already too long but just to give you an idea....

MCA first of all separates the parameters in the cell (or pathway) from the variables in the cell. If one is just studying a pathway then the parameters are usually taken to be the boundary conditions, i.e source and sink pools, the concentration of enzymes and their respective kinetic parameters. For the sake of argument we can treat the expression level of the enzymes to be the 'genotype' of the system. The variables of the pathway are the concentrations of intermediate metabolites and the pathway flux. These are what we consider to be the 'phenotype'. The question that MCA then attempts to answer is how do the enzymes (through their concentrations and kinetic parameters) determine the levels of metabolites and flux through the pathway and how are the metabolite concentrations and flux affected by changes (for example) in enzyme levels. MCA attempts to answer this question by tracing both qualitatively and quantitatively all the cause and effect communication routes through the pathway. In this way we are able to build up a picture of how the pathway is working and what would happen if it were perturbed, say by a mutation.

The moral of the story is, yes go ahead, have a go at simulating the cell, you never know, you might come up with something nobody else has thought of or discovered.

Herbert Sauro e-mail: hrs@aber.ac.uk
Biological Sciences phone: +44 970 622353
Univesity College of Wales

Date: Thu Aug 06, 1992 5:18 am PST
Subject: no sneer

[From: Bruce Nevin (Thu 92086 08:27:26)]

(Avery Andrews (9208060)) --

For the record, my (Wed 92085 12:21:44) post was not intended as sneering, but rather as an apology and an invitation--an apology for seeming to sneer (or whatever it was that ticked Eric off), and an invitation to Eric to help clarify the relationship between the varieties of ALife and PCT.

(Eric Harnden (920805)) --

GAs depend upon Neo-Darwinian negative feedback. That does not necessarily mean that "organisms" evolved through GAs employ negative feedback in their functioning. But of course this does not rule out the possibility of observing the phenomenon of control in these products of ALife research. Perhaps someone should be on the lookout for it, Eric?

If Products of GAs (or any other products of ALife research) exhibit control, it would then be of value to see by what mechanisms they effect control, and whether or not such mechanisms are consistent with PCT.

If these mechanisms are consistent with PCT, then you have a ready-made theory to explain something fundamental (something about achieving consistent ends by variable means in the face of disturbance) which I am reasonably certain no one else in the ALife world has noticed (but I would welcome your proving me wrong). Being in such a position could be a lot of fun for you.

If these mechanisms (which you might find if you look) are not consistent with PCT, then we would greatly appreciate knowing about it.

So we have good reasons for wanting to learn more, and for asking for your help in this.

Bruce bn@bbn.com

Date: Thu Aug 06, 1992 8:54 am PST
Subject: Astro; Die-hards; biochemical systems

[From Bill Powers (920806.0900)]

Avery Andrews (920806) --

>... suppose you want to get from point A to point B in minimum time >(assuming newtonian physics), and energy expenditure & avoiding >obstacles is not an issue. Then you want to go at maximum acceleration >on the first half of your trip, and maximum deceleration on the second >(the time-position plot of the second. half will look like that of a >falling coffee-cup, with time reversed).

I think you're computing output -- trying to determine in advance what the thruster behavior will be, instead of letting the control system vary the thruster output by itself.

Here's my C program for making Astro chase a moving mother ship.

```
#include "stdio.h"
#include "stdlib.h"
#include "string.h"
#include "math.h"
#include "dos.h"


void mother(void);
void astrolev1(void);
void astrolev2(void);

float dt,left,right;
int i;

/* mother ship */
float mv,mx;

/*astro */
float aa, /* astro accel */
      av, /* astro vel */
      ax, /* astro x position */
      pv, /* perception of rel velocity */
      rv, /* reference rel velocity */
      ev, /* rel velocity error */
      px, /* perception of rel x position */
      rx, /* reference rel x position */
      ex, /* rel position error */
      k1, /* velocity control gain */
      k2; /* position control gain */

FILE *data;

void mother()            /* mother ship behavior */
{
 if(mx < left || mx > right) mv = -mv;
 mx += mv;
}

void astrolev1()         /* astro's velocity control */
{
 pv = av - mv;          /* vel perception is difference velocity */
 ev = rv - pv;          /* compute velocity error signal */
 aa = k1*ev;           /* compute acceleration (thruster force) */
 av += aa * dt;         /* integrate accel to get velocity of astro */
 ax += av * dt;         /* integrate vel to get position of astro */
}

void astrolev2()        /* astro's position control */
{
 px = ax - mx;          /* perception of position difference */
 ex = rx - px;          /* position difference error */
 rv = k2*ex;            /* velocity ref prop to position error */
}
```

```
void main()
{
 data = fopen("data","a");
 if(!data) exit(0);
 dt = 0.1;        /* basic time interval */
 mv = 0.2;        /* mother moves at this velocity, back and forth */
 k1 = 1.0;        /* velocity error gain */
 k2 = 3.0;        /* position error gain */
 av = 0.0;        /* start astro at 0 velocity */
 ax = 0.0;        /* start astro at position 0 */
 mx = 10.0;       /* start mother at position 10 */
 rx = 0.0;        /* reference relative position = 0 */

 left = -15.0;    /* limits of mother's motion */
 right = 15.0;

 for(i=0;i<50;++i)
  {
   mother();
   astrolev1();
   astrolev2();
   fprintf(data,"mx:%5.1f  ax:%7.1f aa:%5.1f
                 av:%5.1f\x0d\x0a",mx,ax,aa,av);
  }
 fclose(data);
}
```

Note that Astro's perceptions are astro minus mother, not mother minus astro as I said in my last post.

The first 50 iterations look like this:

```
Mother     Astro          Astro       Astro-Mother
Pos'n      pos'n          accel   velocity (relative vel)

mx: 10.2  ax:    0.0     aa:  0.2      av:  0.0
mx: 10.4  ax:    0.3     aa: 30.8      av:  3.1
mx: 10.6  ax:    0.9     aa: 27.4      av:  5.8
mx: 10.8  ax:    1.7     aa: 23.5      av:  8.2
mx: 11.0  ax:    2.7     aa: 19.3      av: 10.1
mx: 11.2  ax:    3.9     aa: 14.9      av: 11.6
mx: 11.4  ax:    5.2     aa: 10.5      av: 12.7
mx: 11.6  ax:    6.5     aa:  6.3      av: 13.3
mx: 11.8  ax:    7.8     aa:  2.3      av: 13.5
mx: 12.0  ax:    9.2     aa: -1.4      av: 13.4
mx: 12.2  ax:   10.5     aa: -4.7      av: 12.9
mx: 12.4  ax:   11.7     aa: -7.5      av: 12.2
mx: 12.6  ax:   12.8     aa: -9.8      av: 11.2
mx: 12.8  ax:   13.8     aa:-11.6      av: 10.0
mx: 13.0  ax:   14.7     aa:-12.8      av:  8.7
mx: 13.2  ax:   15.4     aa:-13.5      av:  7.4
mx: 13.4  ax:   16.0     aa:-13.8      av:  6.0
mx: 13.6  ax:   16.5     aa:-13.6      av:  4.6
mx: 13.8  ax:   16.8     aa:-13.1      av:  3.3
```

```
mx: 14.0  ax:   17.0     aa:-12.2      av:  2.1
mx: 14.2  ax:   17.1     aa:-11.0      av:  1.0
mx: 14.4  ax:   17.1     aa: -9.6      av:  0.1
mx: 14.6  ax:   17.1     aa: -8.0      av: -0.7
mx: 14.8  ax:   16.9     aa: -6.4      av: -1.4
mx: 15.0  ax:   16.7     aa: -4.8      av: -1.9
mx: 15.2  ax:   16.5     aa: -3.1      av: -2.2
mx: 15.0  ax:   16.3     aa: -2.0      av: -2.4
mx: 14.8  ax:   16.0     aa: -1.7      av: -2.5
mx: 14.6  ax:   15.8     aa: -1.3      av: -2.7
mx: 14.4  ax:   15.5     aa: -1.0      av: -2.8
mx: 14.2  ax:   15.2     aa: -0.7      av: -2.8
mx: 14.0  ax:   14.9     aa: -0.4      av: -2.9
mx: 13.8  ax:   14.6     aa: -0.1      av: -2.9
mx: 13.6  ax:   14.3     aa:  0.2      av: -2.9
mx: 13.4  ax:   14.1     aa:  0.4      av: -2.8
mx: 13.2  ax:   13.8     aa:  0.6      av: -2.7
mx: 13.0  ax:   13.5     aa:  0.8      av: -2.7
mx: 12.8  ax:   13.3     aa:  0.9      av: -2.6
mx: 12.6  ax:   13.0     aa:  1.0      av: -2.5
mx: 12.4  ax:   12.8     aa:  1.0      av: -2.4
mx: 12.2  ax:   12.5     aa:  1.0      av: -2.3
mx: 12.0  ax:   12.3     aa:  1.0      av: -2.2
mx: 11.8  ax:   12.1     aa:  1.0      av: -2.1
mx: 11.6  ax:   11.9     aa:  0.9      av: -2.0
mx: 11.4  ax:   11.7     aa:  0.8      av: -1.9
mx: 11.2  ax:   11.6     aa:  0.7      av: -1.8
mx: 11.0  ax:   11.4     aa:  0.6      av: -1.8
mx: 10.8  ax:   11.2     aa:  0.4      av: -1.7
mx: 10.6  ax:   11.0     aa:  0.3      av: -1.7
mx: 10.4  ax:   10.9     aa:  0.2      av: -1.7
```

I put the data in a file so I could merge it into this post. You could just use a printf (without writing to a file) to put it on the screen.

Note how Astro starts with a burst of acceleration to catch up, then ends up lagging a little behind Mother with small and decreasing acceleration.

I don't mean to take over your project, but I think it will help in learning to do simulations to see a working control system. You may want to do this differently, but seeing one way that works can't do any harm.

RE: Alife

>Diehard PCT-ers believe that living systems have to be made out
>of control systems, but nobody really *knows* this yet.

If control is necessary, then control systems are necessary. So the question comes down (as Rick Marken keeps saying) to whether there's some reason that the behavior of organisms that we observe has to be control behavior or whether it may possibly be some other kind of behavior.

Control behavior is required any time that a regular consequence of actions is seen under circumstances where the organism contributes only one of several influences that produce that consequence. Astro produces a more or less constant distance between itself and Mother, under circumstances where Mother can move independently and Astro can influence its own motions but not Mother's motions. Therefore it exhibits (as I programmed it) control behavior. You wouldn't have to see my program to see that control is involved.

If Astro could sense Mother's motions and convert them into corresponding open-loop thruster actions, some semblance of control could be maintained for a while. But Astro moves according to time integrals of thruster force, so eventually this method of making Astro's movements conform to those of Mother would drift out of calibration and error would increase without limit -- integration errors are cumulative. The only way I know of for Astro to follow Mother INDEFINITELY is to use control loops, as above. I don't think anyone else knows another method, either.

Similar arguments hold for all real behaviors of organisms that I know about. It's impossible for organisms to behave as they do, in the environments where we find them, without being able to control outcomes. It follows that these behaviors must be produced by control systems, the only known kind of system that can control outcomes despite unpredictable disturbances of those outcomes.

So I protest the epithet "Die-hard PCTers." This makes it sound as if we have decided as a matter of faith that all organisms are control systems and are going around trying to prove it. Our arguments with Alife, if they really exist, have to do with their assuming that regular outputs can produce regular outcomes, which is not true in this universe. If Alifers come to realize why control is needed and under what circumstances, and begin devising their models to work in those circumstances, we'll have no argument left. But as long as Alifers or anyone else continues to devise models that work only in special environments free of disturbances, using calculations of infinite precision, I will continue to say that they are not modeling the behavior of organisms.

Oded Maler (920806) --

You forwarded a fascinating post. I know just a little about Biochemical Systems Analysis (Savageau), just enough to know that while feedback has certainly been recognized, control hasn't. The enormous complexity of "paths" will, I think, yield to a control analysis, especially one involving hierarchies of control. What seems to be standing in the way now is a stubborn refusal by biologists to recognize reference signals for what they are. When it's not realized that control systems with reference signals set the concentrations of controlled variables, making those variables independent of the substrate and of disturbances, the whole schmear looks complex beyond redemption. I truly wish we had some biochemists in the CSG, willing to look at these problems from a new point of view.

Best to all, Bill P.

Date: Thu Aug 06, 1992 12:04 pm PST
Subject: Alife and all that

[From Rick Marken (920806)] Oded Maler (920806) says:

> behavior-based
>robotics (a-la Brooks) is just one branch (not the main one)
>in the Alife coalition. The mainstream is more related to things
>such as cellular automata, self-replication, emergent proerties,
>self-organization, complex dynamics and this kind of stuff.

The term "behavior-based" does seem to give away the non-PCT related goals of Brooks' robotics effort. But I am familiar with work done under some of those other rubrics (particularly complex dynamics) and I have found no one in any of those areas who is trying to deal with the fact of control in living systems.

I hate to be a broken record about this (the fact that control theory is about control) but what can I do; nobody in these (and other) fields of the life sciences seems to listen to this point so the inclination is to say it (and try to demonstrate it) over and over; after all, I'm a control system controlling for (amomg other things) the perception that others have heard me.

>I also think that the "crowd" demo is the PCT work which is closer
>in the spirit to some Alife research.

On the surface maybe; but is the spirit of Alife to explain how complex appearing behavior results from the efforts of individual organisms to CONTROL perceptual variables? Do they see twhat is most interesting (and important) about the individuals in "crowd" is the inputs they are trying to produce for themselves, and not the outputs that are so compellingly obvious on the screen?

If so, then, indeed, Alife and PCT are doing the same thing -- and that would be great.

I hope that you (Oded) or Eric can find the time to explain some ALife project in detail and show how it relates to PCT.

Best regards Rick

Date: Thu Aug 06, 1992 1:07 pm PST
Subject: alife...

i find myself in a variety of awkward positions. i will now try to disentagle my body parts...

1) i think i have been mistakenly attributed an expertise which i did not claim. allow me to reiterate the phrase 'interested observer and dilettante.' i am not an alifer. i am not a pcter. i am a modeller in general, and do not willingly adhere to any single paradigm, outside of the personal conviction that cybernetic thinking is right thinking. i do not have, ready to hand, a grab-bag of arguments for any occasion. i am essentially a tinkerer, seeking 'the pattern which connects.' i am drawn here to this group because i find in pct an explanation which seems to work, which is my second criterion for interest in anything. i am, for the same reasons, drawn to studies in neural networks, alife, ecology, semiotics, etc. but i have had too many experiences

(in the lab, in the studio, in my kitchen) of just being wrong, regardless of the strength of my convictions on any given topic. all of these things, then, including pct, remain as models for me. i am in the middle of thinking about these patterns, and am simply not in a position to declare any truth with any authority.

2) despite my initial irritation at his suggestion that alife research is pursuing false goals, i find myself agreeing with the balance of mr nevin's subsequent comments. he has clearly articulated his answer to my own question ('what does he think alife goals are?'), and in so doing has nicely defined the the current distinctions between alife and pct. i would particularly like to express my thanks for his interpolation of my 'thought vs. action' phrase... i was interrupted several times during the writing of that post, and did not make some connections as clearly as i might have done.

3) i think that this distinction is also indicated by mr andrews' comments, pointing out that alife is not a single study, say of robotics. what appears to me to unify the field is the question 'how can organization be seen as a phenomenon?' put differently, how can order arise from disordered elements, without explicitly planning for it? certainly, at this stage in the game there are a lot of simplifying assumptions being made, not the least of which is the lack of certain kinds of disturbance in the medium in which order is to manifest itself. and for certain directions of research, these will continue to be valid simplifications... as in robotics, where much of the activity is oriented toward obtaining particular features of behavior, regardless of the mechanism. this, truly, is simply output modeling. but in other areas i think mr nevin is right... it will be interesting, at some point, to introduce the pct perspective to the alife toolkit, and see if the systems which order themselves exhibit the property of 'control'. however, control is not, in my view at least, synonymous with organization, nor is it the only property of life, however fundamental. hence my contention that alife and pct are getting along quite nicely without each other at the moment. (to the extent that they are in fact getting along without each other. their distinctions can be overdrawn, i think.) when the processes of organization are well enough understood that they can be adequately predicted and interpreted, then further modification to deal with problems of 'consistency in disturbance- prone environments' will be appropriate. but let's play with one variable at a time, shall we?

4) thinking about this has just brought a question to mind, the answer to which might significantly increase my understanding of control theory, and close the circle on this argument...

if that which is perceived (the effect of action on the environment) is perceived without distortion (no noise, no disturbance, no delay), and if the environmental response to action is at the very least deterministic (possibly linear, certainly consistent, maybe even just transparent), then isn't the perception of the outcome equivalent to the perception of the action? it occurs to me that this is the central simplifying assumption that leads to the criticism that alife is not in fact modeling living systems (living systems, by pct definition, controlling outcome and not action, as oft repeated). if so, i hold to my position that the question most important to alife ('how can anything organize') needs to be thoroughly addressed before questions of a different order ('how can anything remain stable') are put forward.

5) so i will not claim that living things are not control systems. nor will i claim that alife is engaged in producing things that are fundamentally like living things. neither will i claim, in fact, that any model which focuses on a given phenomenon can produce anything which is fundamentally like living things. different features of interest require, for the moment, different experimental approaches. and if the bug walks, or the swarm gathers, or the predator kills, or the eye tracks, or the perceiver controls... then the model is informative.

-----------< Cognitive Dissonance is a 20th Century Art Form >-----------
Eric Harnden (Ronin)
<HARNDEN@AUVM.BITNET> or <HARNDEN@AMERICAN.EDU>
The American University Physics Dept.
4400 Mass. Ave. NW, Washington, DC, 20016-8058
(202) 885-2748
---------------------< Join the Cognitive Dissidents >-------------------

Date: Thu Aug 06, 1992 7:08 pm PST
Subject: astro & alife

My astro code works pretty much like Bill's, other than the square root business, which could be spliced in by replacing:

```
  ex = rx - px;        /* position difference error */
```

with

```
  ex = root(rx - px)
```

where root(x) is defined as

```
  root(x) = sign(x)*sqrt(abs(x))
```

and sign(x) = 1 if x >= 0, -1 otherwise.

So its just a twist in the perceptual function that lets astro close in faster on mother without overshooting. If there were self-reproducing robots whizzing around in space, natural selection could cause them to acquire this function, since its utility follows from fixed properties of the environment.

I'm beginning to acquire some confidence that astro's reorganization really does confine itself to a region for k2 (in Bill's C code), fluctuating between .2 and .9, where about .55 seems to be about optimal. I think the fluctations are due to the intervals between choice of new delta being too short, so that differences between current and old average error are too much influenced by accidents of how mother is moving around. But I think I'll have to redo it properly in C (it's currently Turbo Prolog (!!! ... because of the nice IDE)) to get it to run at a reasonable rate.

From an empirical point of view, does anyone know anything about the distance - velocity functions used by birds when landing? It might be interesting to see how living systems actually tackled the problems posed by inertia in a low-friction environment.

As for Alife, etc: Many of these systems (and also Chapman & Agre's video-game playing programs) do model significant aspects of keeping oneself alive (that's why video games are fun). So either they are in fact full of control systems, perhaps to a greater extent than their creators realize, or they are leaving out aspects of reality for which control systems are essential. Either way they provide lots of stuff for people to do, either in the way of improving our understanding of how they work, or in making them more lifelike, or both. This is pretty much what Greg Williams has been saying, I think, and of course his & Pat's NSCK program is a fabulous resources for helping people to get on with it.

Avery.Andrews@anu.edu.au

Date: Fri Aug 07, 1992 5:23 am PST
Subject: Re: Alife and all that. Specifications

[From Oded Maler 920807]

[Rick yesterday:]

>I hope that you (Oded) or Eric can find the time to explain some ALife project
>in detail and show how it relates to PCT.

Sorry to disappoint you but like MR HARNDEN (whose general attitude seems to be like mine) I'm not an active member of the alife community (except for a paper on my worm rejected from the first ECAL, which I attended, and except for my bad habit of tracking "trendy" branches of "science"). You might get idea of their work from looking at their proceedings (usually edited by Langton, Addison-Wesley) - I'm sure you'll find non-references to PCT almost in every page :-)

Since you mentioned complex dynamics, I recall hearing in Aix a talk by Kugler about changing observables and all that. Do you have any opinion (surprising or not, it doesn't matter) on this stuff?

Concerning the Astro discussion, since I'm interested in an area which might be called "specification and verification of non-living syatems" I would like to ask what do you think will be an appropriate specification of the properties of the Mother-Astro systems. That is, I'm looking for something of the form:

Given that:

1) some properties of engines (gains), sensors (delays)
2) some restriction on the patterns of movement of Mother (e.g., it will not switch from positive speed of light to negative one too often)
3) some limitation on the noise (maybe a property of sensors and actuators)

Then:
e.g.,:
1) there will be some moment t where the distance remains smaller than d

2) for every time t when distance is bigger then d there is t'>t
   where distance is smaller then d.
3) for every interval the integral/average of distance is smaller than d.

I'm interested in developing methods for giving formal proofs for such properties (at this stage mathematical control theory seems to be sufficient but when discrete events appear, e.g., Mother is turning into a Monster (like in PacMan) it becomes more complicated.)

Thanks --Oded

Date: Fri Aug 07, 1992 5:47 am PST
Subject: pct<->alife, or why should meat want?

from: eric harnden(920807)

sorry i forgot the time stamp on my last post.

i had a conversation with my research partner this morning that helped me to make the relation between alife and pct a little crisper, at least to myself.

as i understand it, any control system requires a reference value, or goal state, in order to operate. the proverbial thermostat is programmed externally to control air temperature for a certain value. a windup toy can be seen as controlling for equilibrium (don't beat on me for that one... it's only an illustration). an ant, while not a windup toy, still controls for some condition... let's say lack of hunger, for instance. but the ant is not, in my view, externally programmed. its reference values arise from the particular genetic conditions which effect the organization of its constituent chemicals, modified somewhat by immediate environmental input. they are programmed, less as in the case of the thermostat where they must be externally supplied, than as in the case of the windup toy, where they arise as a consequence of structure. the question that is of interest to me, and i think that makes alife relevant, is why matter should organize itself in this way at all... and, taken further, what properties of ordering that characterize lower-order organization may also shape higher-order 'behaviors' (the act of organization being seen as a kind of behavior). given that the difference between crystalline and organic growth is that the organism perceives and controls, how does this distinction (among other properties of life and mind) arise? purposive behavior (the essence of 'mind') requires that purpose exist. if not provided exogenously, it must be immanent in the system, and behavior emergent as a result (a by-product, i am told). am i being either naive or unecessarily elliptical here? or does this go any farther toward clarifying a position?

-----------< Cognitive Dissonance is a 20th Century Art Form >-----------
Eric Harnden (Ronin)
<HARNDEN@AUVM.BITNET> or <HARNDEN@AMERICAN.EDU>
The American University Physics Dept.
4400 Mass. Ave. NW, Washington, DC, 20016-8058
(202) 885-2748
---------------------< Join the Cognitive Dissidents >-------------------

Date: Fri Aug 07, 1992 6:13 am PST

Subject: how can anything organize?

[From: Bruce Nevin (Fri 920807 08:07:45)]

(Eric Harnden (Thu, 6 Aug 1992 14:19:35 EDT)) ---

>what appears
>to me to unify the field is the question 'how can organization be seen as
>a phenomenon?' put differently, how can order arise from disordered elements,
>without explicitly planning for it?

>the question most important to alife [is] 'how can anything organize'

Just before everybody disappeared for the annual Durango bash, I made some proposals about the relation between control systems of different orders of structural complexity. Let me restate these in this new context as follows:

Structure of order n at least sometimes "emerges from" (is a colligative property of) the behavioral outputs of control systems of order n-1, each controlling its individual perceptions. The "gather" program illustrates this nicely.

Control systems may perceive aspects of such structure, and may come to control those perceptions. (Is there anything interesting going on? There's something. I wonder why they're all gathering around her? Must be some reason to pay attention to her. I'll think I'll try to get as close as I can.) They may create and may come to depend upon conventions and institutions that reflect these perceptions of structure (e.g. build amphitheaters).

Control systems of order n (e.g. multicellular) can evolve from structures emergent from control by control systems of order n-1 (e.g. cells). There seems to be a vast evolutionary "distance" between e.g. fungae vs. bacterial colonies (it seems likely that fungae are control systems?) on the one hand and neural ECSs vs. neurons on the other hand. The differences seem to reside in interlocking cellular specializations, each constituting part of an environment on whose reliable structure the other cells depend, and for the maintenance of which they each in their several ways control.

My focus previously was on "communication" downward from order n to order n-1: reorganization, I suggested, is the result of CSs of order n-1 controlling to reduce environmental "toxicity", while concurrently CSs of order n, so long as they continue to experience conflict (and consequent chronic error), produce intrasomal effects that happen to be "toxic" to their constituent elements of order n-1, more or less localized around the order-n CSs in conflict and needing reorganization.

The focus now is on "communication" upward from order n-1 to order n, and on the question of how the later may first evolve from and then be continuously recreated and sustained by the former.

I know that others here (Bill, Greg, Gary, Francis Heylighen, others) have been involved in the relation between control theory and evolution and know a heck of a lot more about it than I. Do these hunches seem feasible?

Bruce bn@bbn.com

Date: Fri Aug 07, 1992 6:13 am PST
Subject: hypercard stacks

[From: Bruce Nevin (Fri 920807 09:29:05)]

Martin,

Some time back you offered your hypercard stacks. Mirabile dictu, I had just been thinking about them a day or two earlier, supposing it would be wonderful to have access to them, wondering if you would be willing to make them available--and wishing I had a mac. When I do, you may be sure I will apply to the server for a copy. Meantime, I appreciate your (continued) generosity.

(Hm--I wonder if it is possible to convert hypercard stacks to any of the hypercard lookalike programs under DOS . . . anybody know anything? I'll make some inquiries.)

Bruce bn@bbn.com

Date: Fri Aug 07, 1992 7:48 am PST
Subject: Alife models; Astro with sqrt

[From Bill Powers (920807.0700)]

Eric Harnden (920806) --

RE: Alife research:

> ... certainly, at this stage in the game there are a lot of simplifying >assumptions being made, not the least of which is the lack of certain kinds of >disturbance in the medium in which order is to manifest itself. and for certain >directions of research, these will continue to be valid simplifications...

>... if that which is perceived (the effect of action on the environment) is >perceived without distortion (no noise, no disturbance, no delay), and if the >environmental response to action is at the very least deterministic (possibly >linear, certainly consistent, maybe even just transparent),then isn't the >perception of the outcome equivalent to the perception of the action?

Well, you've defined the problem quite clearly. Simplifying assumptions are always necessary in modeling; if they couldn't be used, Schroedinger's equation would be the basis of all models of everything (and even it would be an approximation). Some simplifying assumptions, however, are more critical than others -- especially if they happen to be false to fact in an important way.

If perception could be equated to the objective physical state of the environment, and if regular effects on the environment could be traced back to regular outputs of the nervous system, cause-effect models such as the stimulus-response model would work perfectly well. That is why they were invented -- because everyone believed in exactly the simplifying assumptions you propose. Everything I have seen in AI or AL tells me that these assumptions are still being made. Without them, there's hardly any proposition about living systems that would work, those propositions that somehow entail closed-loop control processes aside.

In fact neither basic assumption is correct; neither the one concerning input nor the one concerning output. Perceptions do not correspond accurately (and often do not correspond at all) to objective events or states, and outcomes of action are almost never regular functions of the outputs of the nervous system. No theory is needed to prove either statement. All that's necessary is to stop making these assumptions and look at real behavior.

The easiest proposition to falsify is the one saying " ...the environmental response to action is at the very least deterministic (possibly linear, certainly consistent, maybe even just transparent) ..". The environmental response to organismic actions can often be modeled as linear, but it can't often be modeled correctly as deterministic or consistent, and never can be modeled as transparent. Action does not determine outcomes, nor are the outcomes of a given action consistent. A given action can easily have opposite effects on an outcome from one trial to the next. And the connection between all actions and their consequences is mediated by properties -- usually variable properties -- of the external world.

One has to search for very special circumstances to find cases in which a given action always has the same outcome, or in which a given outcome is repeatable by repeating the same action. In fact, one has to set up a laboratory experiment in which all normal variations in the environment are suppressed, in which most normal causal links in the environment are missing and the remaining ones remain precisely repeatable, and in which the organism itself always begins the experiment in the same state. The purpose of setting up such laboratory conditions is precisely TO MAKE THE ASSUMPTIONS APPEAR TRUE. And even when such conditions are established, variability still appears. All attempts to force organisms to behave according to these "simplifying" assumptions have failed. Yet those who purport to be studying the behavior of living systems have stubbornly insisted that the assumptions must be correct.

If they aren't correct, all theories that assume their correctness are simply wrong. They aren't just a little off because of making some simplifying assumptions. They're qualitatively wrong and their predictions are quantitatively as well as qualitatively wrong. Real organisms in real environments would not behave as models that make these assumptions behave.

If you'll examine any behavior without making these simplifying assumptions, you'll see that that there's a clear answer to your question, "isn't the perception of the outcome equivalent to the perception of the action?" The answer is "No." It's OBVIOUSLY "no," once you look carefully at the way any behavior is produced.

The primary way in which people have made the answer seem to be "yes" is by naming behaviors not in terms of actions but in terms of outcomes. When you see someone hammering a nail into a board, you're seeing the outcome of a series of efforts applied to a hammer, and you're naming the action by its outcome: the hammer hits the nail and the nail sinks into the wood. Neither of those processes is the action of the carpenter; they are environmental processes. If the carpenter generated exactly the same motor command signals over and over, the hammer would miss the nail most of the time. When you say a person is steering a car down a road, you're referring to the fact that the steering wheel somehow turns in just the way needed to keep the car on the road; you're naming environmental processes, not the actions by which the driver is making those processes occur. In the case of steering a car, it's clear that the efforts applied to the steering wheel are almost unrelated to the car's observed path or even the position of the steering wheel. If the efforts were repeated exactly, the car would be in the ditch in a jiffy.

When disturbances are omitted from a behavioral model, this is not just a simplifying assumption. Omitting disturbances conceals the fundamental feature of a control system. A model designed specifically to work in an environment where outcomes and actions are the same thing will not just work a little worse when disturbances are introduced: it will not work at all.

In normal environments, disturbances have at least as much effect on outcomes as actions do; they often require reversing an action in order that the same outcome repeat. In most behaviors, the PRIMARY reason that continued action is needed is that disturbances need to be counteracted. If you were driving a mechanically perfect car down a perfectly straight level smooth road in absolutely still air, you could just aim the car in the right direction with a telescopic sight, clamp the steering wheel, and then read a book for the rest of the trip. In most real environments, outcomes are not simply proportional to actions; they involve time integrals and other processes extending through time. Even small errors in action are magnified to create large differences in outcomes.

Even without disturbances, real outcomes of real actions are variable; they are hypersensitive to initial conditions. Models that try to do without negative feedback usually have to employ calculations of extreme precision in order to continue working for even modest lengths of time; if required to act continuously, each action starting where the previous one left off, the calculations will depart further and further from reality, just as Lorenz found in his weather models. In one mode of my arm model, the Little Man points to a target that jumps randomly about in space every 3/4 second. I can leave it running overnight (and have done so), and in the morning it's behaving just as it was when I went to bed, despite the fact that the integrations are done in 16-bit integer arithmetic. If you tried that with any of the open-loop models being proposed, even using 80-bit floating point arithmetic, you'd probably awaken to find an overflow error on the screen and the Little Man with his arm wrapped around his neck.

I strongly reject the eclectic approach you imply when you say

>different features of interest require, for the moment, different experimental >approaches. and if the bug walks, or the swarm gathers, or the predator kills, >or the eye tracks, or the perceiver controls... then the model is informative.

This approach treats all theories as if they were indistinguishable. In psychology perhaps there's a good reason for this: all theories are equally bad, so it doesn't matter much which one you use -- you'll still get correlations of 0.5 or so. But in modeling behavior, one can easily choose between models that predict behavior in the real world and those that predict it only in a world based on assumptions known to be false. Control theory clearly belongs with theories that allow for perception to be a construction rather than a veridical report on reality, and for action to be only one influence on outcomes among many. I don't know exactly how many theories fall in this class, but my guess is 1. All the other theories are based on a known falsehood; I find no trouble with rejecting them simply for that reason.

If you can think of an example of a real behavior in which outcome is a quantitatively regular function of action, I'd like to hear about it.

Avery Andrews (920806) --

I tried the square-root substitution. With dt = 0.1, k1 = 3, and k2 = 5, the behavior goes like this:

```
Mother      Astro         Astro      Astro-Mother
Pos'n       Pos'n         Accel         vel

mx: 10.2  ax:    0.0     aa:  1.0     av:  0.1
mx: 10.4  ax:    0.5     aa: 48.4     av:  4.9
mx: 10.6  ax:    1.2     aa: 23.5     av:  7.3
mx: 10.8  ax:    2.1     aa: 10.5     av:  8.3
mx: 11.0  ax:    2.9     aa:  3.7     av:  8.7
mx: 11.2  ax:    3.8     aa:  0.1     av:  8.7
mx: 11.4  ax:    4.7     aa: -1.8     av:  8.5
mx: 11.6  ax:    5.5     aa: -2.7     av:  8.3
mx: 11.8  ax:    6.3     aa: -3.2     av:  7.9
mx: 12.0  ax:    7.0     aa: -3.5     av:  7.6
mx: 12.2  ax:    7.8     aa: -3.6     av:  7.2
mx: 12.4  ax:    8.5     aa: -3.6     av:  6.9
mx: 12.6  ax:    9.1     aa: -3.6     av:  6.5
mx: 12.8  ax:    9.7     aa: -3.5     av:  6.2
mx: 13.0  ax:   10.3     aa: -3.5     av:  5.8
mx: 13.2  ax:   10.8     aa: -3.4     av:  5.5
mx: 13.4  ax:   11.4     aa: -3.4     av:  5.1
mx: 13.6  ax:   11.8     aa: -3.3     av:  4.8
mx: 13.8  ax:   12.3     aa: -3.2     av:  4.5
mx: 14.0  ax:   12.7     aa: -3.0     av:  4.2
mx: 14.2  ax:   13.1     aa: -2.9     av:  3.9
mx: 14.4  ax:   13.5     aa: -2.8     av:  3.6
mx: 14.6  ax:   13.8     aa: -2.6     av:  3.4
mx: 14.8  ax:   14.1     aa: -2.4     av:  3.1
mx: 15.0  ax:   14.4     aa: -2.2     av:  2.9
mx: 15.2  ax:   14.7     aa: -1.9     av:  2.7
mx: 15.0  ax:   14.9     aa: -3.7     av:  2.3
mx: 14.8  ax:   15.1     aa: -8.2     av:  1.5
mx: 14.6  ax:   15.1     aa:-16.3     av: -0.1
mx: 14.4  ax:   14.9     aa:-10.6     av: -1.2
mx: 14.2  ax:   14.8     aa: -6.2     av: -1.8
```

```
mx: 14.0  ax:   14.5     aa: -3.3     av: -2.1
mx: 13.8  ax:   14.3     aa: -1.5     av: -2.3
mx: 13.6  ax:   14.1     aa: -0.5     av: -2.3
mx: 13.4  ax:   13.9     aa:  0.1     av: -2.3
mx: 13.2  ax:   13.6     aa:  0.4     av: -2.3
mx: 13.0  ax:   13.4     aa:  0.5     av: -2.2
mx: 12.8  ax:   13.2     aa:  0.5     av: -2.2
mx: 12.6  ax:   13.0     aa:  0.4     av: -2.1
mx: 12.4  ax:   12.8     aa:  0.4     av: -2.1
mx: 12.2  ax:   12.6     aa:  0.3     av: -2.1
mx: 12.0  ax:   12.4     aa:  0.2     av: -2.0
mx: 11.8  ax:   12.2     aa:  0.2     av: -2.0
mx: 11.6  ax:   12.0     aa:  0.1     av: -2.0
mx: 11.4  ax:   11.8     aa:  0.1     av: -2.0
mx: 11.2  ax:   11.6     aa:  0.0     av: -2.0
mx: 11.0  ax:   11.4     aa:  0.0     av: -2.0
mx: 10.8  ax:   11.2     aa:  0.0     av: -2.0
mx: 10.6  ax:   11.0     aa:  0.0     av: -2.0
mx: 10.4  ax:   10.8     aa: -0.0     av: -2.0
mx: 10.2  ax:   10.6     aa: -0.0     av: -2.0
mx: 10.0  ax:   10.4     aa: -0.0     av: -2.0
mx:  9.8  ax:   10.2     aa: -0.0     av: -2.0
mx:  9.6  ax:   10.0     aa: -0.0     av: -2.0
mx:  9.4  ax:    9.8     aa: -0.0     av: -2.0
```

So a nonlinear comparator doesn't make much difference. If you double the loop gain the behavior near zero error will be about the same (sqrt(x) expands to approximately x/2 for small x); the gain will be lower for larger errors, but if the minimum loop gain is high errors will never become large.

One thing to watch out for when you're reorganizing this or any other system that's simulated on a computer. When the loop gain gets high enough you can start getting computational artifacts due to the discrete nature of the variables. The value of dt should always be made small enough so it takes several iterations -- five is a nice number -- to make the fastest change in a variable. Sometimes computational artifacts can be confused with real instability of the system.

When you say that k2 fluctuates around .55, I assume that you're using a dt of 1. If you use a smaller dt, you may find that k2 can become a lot larger. You may be looking at a computational artifact (but the fact that reorganization can compensate for it is interesting anyway).

Inertia really isn't much of a problem for control systems in a frictionless environment. If Mother and Astro were in orbit, on the other hand, the problem might become more complex, because you don't move in the direction of thrust relative to another object in the same orbit. The most difficult environments, in terms of physical properties, are those involving slip-stick friction, thresholds, and limits.

>As for Alife, etc: Many of these systems (and also Chapman & Agre's
>video-game playing programs) do model significant aspects of keeping
>oneself alive (that's why video games are fun). So either they are in
>fact full of control systems, perhaps to a greater extent than their

>creators realize, or they are leaving out aspects of reality for which >control systems are essential. Either way they provide lots of stuff >for people to do, either in the way of improving our understanding of >how they work, or in making them more lifelike, or both.

My problem with this generous view is that it's hard to know when such models tell us something about life and when they just inform us about the consequences of playing an arbitrary game. I'm all for video games, but before I can accept any of them as models of living systems, I want to see some explanation of why the rules are relevant to something about living systems.

In a lot of models (like Pengi) the critical part that makes them work isn't even in the model; it's in the modeler.

Best to all, Bill P.

Date: Fri Aug 07, 1992 11:49 am PST
Subject: Gopher Update

The Gopher server on biome has been updated, but is not fully functional. You can now telnet to biome and login as gopher; if your system does not set the terminal type for you, gopher will assume a vt100.

The mail facility has been disabled, and some files cause strange behaviour. I'll fix this when I can, but it won't be for a while.

Bill Silvert

Date: Fri Aug 07, 1992 2:53 pm PST
Subject: ALife

[From Rick Marken (920807)]

First, a big THANK YOU to Bill Powers for that beautiful post (920807) about why it is sometimes not such a hot idea to make "simplifying assumptions" -- it's not such a hot idea when you know those assumptions are wrong. I don't know how many times we will have to repeat (and repeatedly show) that the fundemental assumption of the life sciences is WRONG and that that wrongness is not a trivial mistake. The basic simplifying assumption made by ALifers (and the rest of the behavioral and life sciences) is wrong in just the way necessary to justify continued reliance on cause-effect models of behavior (and cause-effect based methods for testing the match of the model to behavior -- when that is done, which in fields like ALife seems to be rare.) Once you understand that organisms CONTROL then you have to eventually realize that they can do this only because they are organized as control systems -- that control their own perceptual experiences. And the difference between a control system and a cause-effect system is NOT small; it's the whole enchilada when it comes to understanding the nature of the behavior of living systems.

Avery Andrews says:

>As for Alife, etc: Many of these systems (and also Chapman & Agre's
>video-game playing programs) do model significant aspects of keeping oneself

>alive (that's why video games are fun). So either they are in fact full
>of control systems, perhaps to a greater extent than their creators
>realize, or they are leaving out aspects of reality for which control
>systems are essential. Either way they provide lots of stuff for people
>to do, either in the way of improving our understanding of how they
>work, or in making them more lifelike, or both.

Again, I should point out that there are many people running around building what they see as S-R systems that are actually control systems. The "Braitenberg Vehicles" (sp?) are a nice example; also a flocking bird model. These models (when they work) work because they deal with continuous variables, are closed loop with the sign of the feedback negative, and they have the proper gain and dynamics (slowing) to keep them stable. There is no explicit reference -- so the perceptual signal is kept equal to 0 and the external correlate of the perceptual signal is kept at the value that corresponds to 0 perceptual signal. In fact, they are control systems!

The fact that these projects exist with their developers calling them S-R machines suggests to me that your claim that they can "improve our understanding" is, as Bill says, rather generous. These people have absolutely no idea that they are dealing with control systems and they would probably rail at the suggestion that their machines are controlling perceptual variables. The net result of seeing these as S-R machines is to approach the process of "improving" their behavior as a problem of finding more effective means of generating OUTPUT. Of course, this just leads to dead ends or the design of systems that live in worlds of simplifying assumptions. So rather than improving our understanding, I would argue that such efforts actually MASK our ability to progress in our understanding of behavior -- by making it SEEM like behavior can be generated by a cause-effect or output generation model. I agree that the people making these models are VERY CLOSE to a useful approach to understanding behavior; but, then, so was Skinner. Like Skinner, I don't think that these folks really want to take the leap (apparently) to the realization that closed loop negative feedback organizations control PERCEPTION -- so their models end up acting more like camoflage than beacon.

Oded Maler (920807) says:

re :ALife

>You might get idea of their work from looking at their
>proceedings (usually edited by Langton, Addison-Wesley) - I'm sure
>you'll find non-references to PCT almost in every page :-)

Actually, I'm on the Santa Fe Institute's mailing list and saw the proceedings (I think). So I do know a bit about what ALife is about (and, indeed, they have it just as wrong as we imagined).

>Since you mentioned complex dynamics, I recall hearing in Aix a talk
>by Kugler about changing observables and all that. Do you have any
>opinion (surprising or not, it doesn'r matter) on this stuff?

I think I wrote a couple papers related to this stuff. Both are in my Mind Readings book (remember that one, folks?) in the section on coordination. The stuff is a bit complex (mathematically) for me; I think I understand the basic goal of the "complex systems" people -- but they might disagree (and do) about my understanding. I think the idea is that when you have systems with lots of degrees of freedom all varying simultaneously, certain functions of all these degrees of freedom will stabilize -- I think this is basically what an "attractor" is. The attractor is their idea of the goal of the system.

I have two beefs with the complex systems people:

1) it am not convinced (and they don't claim this is true either) that these systems can reach the attractor state in the presense of continuous disturbances and

2) they don't say what set's the parameters of the system so that one particular attractor point (rather than another) is reached -- more important, they don't say why a particular attractor state is achieved (rather than another).

I guess the bottom line for me on complex systems models -- I don't like models that I can't explain to my kids (well, they're not really "kids" any more). I can explain the basic idea of how you maintain your balance in control theory terms. I think I can sort of explain it in "complex systems" terms -- but when I do it in a professional paper I'm always told I don't have it right. So maybe I just don't like theories that I can't understand.

Best regards Rick

Date: Fri Aug 07, 1992 3:48 pm PST
Subject: Re: Alife models

(from penni sibun 920807) [From Bill Powers (920807.0700)]

> As for Alife, etc: Many of these systems (and also Chapman & Agre's
> video-game playing programs) do model significant aspects of keeping
> oneself alive (that's why video games are fun). So either they are in
> fact full of control systems, perhaps to a greater extent than their
> creators realize, or they are leaving out aspects of reality for which
> control systems are essential. Either way they provide lots of stuff
> for people to do, either in the way of improving our understanding of
> how they work, or in making them more lifelike, or both.

My problem with this generous view is that it's hard to know when such models tell us something about life and when they just inform us about the consequences of playing an arbitrary game. I'm all for video games, but before I can accept any of them as models of living systems, I want to see some explanation of why the rules are relevant to something about living systems.

well, i've seen the crowd demo (thanks to avery) and i've seen sonja (pengi's daughter). quite frankly, you know, they're both just demos, and one can look at them, or even play with them, and assume they're video games that aren't trying to demonstrate any particular theory of activity. if you read the accompanying text and/or the code, you may or may not speculate about whether or how the demos support the models.

(don't get me wrong--i think both demos are cool.)

In a lot of models (like Pengi) the critical part that makes them work
isn't even in the model; it's in the modeler.

hm. maybe you're conflating something here. pengi isn't a model, it's an implementation (a demo) of a model. building pengi or crowd (or astro!) required some programmer to figure out what code to write to make the right thing happen. seems in both cases, the coder (the modeler) is firmly in the loop, and indeed critical. but that's a property of the demo, not of the model or theory the demo is meant to demonstrate.

[From Rick Marken (920807)]

>As for Alife, etc: [same excerpt]

Again, I should point out that there are many people running around building
what they see as S-R systems that are actually control systems. The

i plead greater ignorance to alife than anyone else: is this what the alife people claim they are building? it is emphatically not what agre, chapman, brooks and co claim they are building. for these folks and others in ai, for gibsonian psychologists, for ethnomethodologists, etc., agents and their worlds are *not separable*. how can you have a respond to b's stimulus when a and b are the same thing? that's actually not a terribly useful simplification of the argument; try this one: how can you individuate a's-stimulus and b's-response when you can't individuate a and b?

cheers. --penni

Date: Fri Aug 07, 1992 5:06 pm PST
Subject: astro & philosophy

[Avery Andrews (920808)] Bill Powers (920807.0700)

Square roots -

Well, the two methods seemed to produce rather different-looking results when I originally tried them. But of course all sorts of things can go wrong with demos... When I get the C version going (don't hold your breath), I'll include a switch so people can try both.

Generosity -

>My problem with this generous view is that it's hard to know when such
>models tell us something about life and when they just inform us about the
>consequences of playing an arbitrary game.

I think this just calls for common-sense judgement. And surely nobody thinks of video-game playing programs as complete models of living systems. But I would insist that the fact that the games are fun is real evidence that they model significant aspects of keeping yourself alive (since mammals evolved a sense of fun as a mechanism for tuning survival skills). And of course the point of being generous at the outset is so that you can make a real impression later.

The recent discussions by you and Rick do give me a more vivid sense than I had before that these gameplayers operate in a low-disturbance environment (Agre sees the world as a primarily stable and benign place, rather than the flaky and often downright hostile place it really is. The perspective of a programmer as opposed to a real engineer, perhaps?). But there may be a matter of levels here -- maybe at certain higher levels the world is relatively stable and cooperative, or to be more precise, looks that way, since 600 m.y.'s worth of lower level control systems are keeping the flakiness under control. But I think it is us who will have to do the work of souping these models up so that they can cope with more in the way of disturbances -- it does no good to tell them that they ought to (we have seen this method fail to work several times in the last six months).

Maybe we could come up with something that could be licensed to After Dark...

Rick Marken (920807)

Well, I only think that most of these people aren't idiots, & that if you take their systems apart and put them back together with radical improvements, they will take notice. But you've been at this longer than I have -- maybe you're right after all.

And we ought to remember that a mere ten years ago `everybody' thought that you needed detailed plans to do anything in the world. By demolishing this idea, the interactive AI people have created a much more favorable environment for PCT than used to exist.

Avery.Andrews@anu.edu.au

Date: Sat, 8 Aug 1992 09:12:05 -0300
Subject: hypercard stacks

>From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
>Subject: hypercard stacks
>
>[From: Bruce Nevin (Fri 920807 09:29:05)]
>
>Martin,
>
>Some time back you offered your hypercard stacks. Mirabile dictu, I had
>just been thinking about them a day or two earlier, supposing it would
>be wonderful to have access to them, wondering if you would be willing
>to make them available--and wishing I had a mac. When I do, you may be
>sure I will apply to the server for a copy. Meantime, I appreciate your
>(continued) generosity.

Martin kindly sent these to me and they are available from the BIOME server. Just a small sampling, since the volume is very large.

Bill Silvert

Date: Sat, 8 Aug 1992 21:08:48 PDT
Subject: Models and demos -- what sells

[From Rick Marken (920808)]

The recent post by penni sibun got me to thinking about the problem of making impressive demos of control phenomena.

penni sibun (920807) says:

>well, i've seen the crowd demo (thanks to avery) and i've seen sonja
>(pengi's daughter). quite frankly, you know, they're both just demos,
>and one can look at them, or even play with them, and assume they're
>video games that aren't trying to demonstrate any particular theory of
>activity.

>(don't get me wrong--i think both deontrol models and demos is the same

I believe that one of the problems confronted by those of us who are trying to "sell" PCT with models and demos is the same as the problem we confront when trying to point out to psychologists that there is a phenomenon (called "control") that is going on in front of their eyes that they have not taken into consideration in their attempts to understand mind and behavior. The problem is that the disturbances, constraints and calibration problems that make control necessary and obvious are simply invisible. When you point your finger at a target, the pointing just seems to happen; the fact that you can repeat this pointing with great precision seems completely unimpressive. You just point at the target again and again. Disturbances (such as changes in your orientation with repect to gravity), constraints (such as the fixed length of the segments of the arm) and calibration problems (like the fact that a neural signal never produces exactly the same amount of muscle tension) go completely unnoticed. When disturbances are visible (such as movements of the target) they look like stimuli guiding the response. It is, thus, very easy for those who want to, to ignore control.

The "little man" demo is a very impressive piece of graphics programming. But what is really amazing about the little dude is that he (or she?) CONTROLS. But this aspect of the demo is very hard to see. In fact, there are no disturbances in this demo other than target movements. It might help if you could "tug" on the arm and see that this had little effect on the accuracy of pointing. But basically, the problem with demonstrating models of control is the same as the problem of seeing control in normally occuring behavior; what is most amazing about control is what you can't see. And you can't see the amazing aspect of control (disturbance resistance, constraint satisfaction, and calibration compensation) because control itself prevents these things from having any noticable effect. So it is the fault of control itself that the process of control is invisible.

In order to see control, you must be the agent of disturbance; you must be able do something that you know should have an effect on a variable if it were not under control. If you think a person is controlling the poition of a limb then you can literally "push" on the limb to see if the push has the expected effect (movement of the limb). This "test" must be done carefully -- not too much disturbance (control systems have limits to the amount of output they can produce) and an appreciation that control of some variables occurs more slowly than others (so the disturbance may seem to have an an effect but will be slowly cancelled if there is control). My "mind reading" demo is an attempt to show that it is difficult to see control (the movements of all 5 numbers around a computer screen appear to be behaviors of the subject) but that it is absolutely necessary to know what is being controlled in order to know what a person is doing.

I don't know if there is any real dramatic way to show control; we keep trying but we obviously haven't found a real "grabber" that would get psychologists to throw up their hands en masse an cry "oy vay, I've been missing the point for my whole career; people don't respond to stimuli or generate outputs -- they control. Now I have to abandon all my work and start studying control. Damn, how did I miss that -- I guess that guy Powers was wasn't just a stubborn, contrarian radical outsider after all."

Ah. Someday.

penni goes on:

>is this [S-R modeling (rm)] what the
>alife people claim they are building? it is emphatically not what
>agre, chapman, brooks and co claim they are building. for these folks
>and others in ai, for gibsonian psychologists, for
>ethnomethodologists, etc., agents and their worlds are *not
>separable*. how can you have a respond to b's stimulus when a and b
>are the same thing? that's actually not a terribly useful
>simplification of the argument; try this one: how can you individuate
>a's-stimulus and b's-response when you can't individuate a and b?

My comments about people building control systems and calling them SR devices were based on an article I read in a computer graphics magazine where little "bugs" moved around the 2-D surface of the screen, avoiding obstacles and, maybe, chasing other bugs. They were described as SR devices because they were -- then sensed some input (S) than was converted into a force (R) than moved the bug. So the bug is SR but it was in a closed loop SITUATION; the output continuously influenced the input. So there was an RS connection too. The integration that produced movement provided dynamic stability and the effect of force on input was such that increases in force led to a reduced tendency of the input to generate output -- that is, the sense of the feedback from output to input was negative. So these bugs were control systems. But the authors showed NO EVIDENCE of knowing this -- and they carried on about the wonders of their SR model. It was wonderful -- but it was a control system. If they had known that, their research could have taken off in an exciting direction. I'm afraid instead they will end up just trying to develop more and more elaborate methods of generating outputs in order to make the behavior of the bugs even more complex.

I know that many scientists are aware of the fact that there is a relationship between output and input (ie feedback). They just don't know what this means. I maintain that, whatever they call their models, they are thinking of them in SR terms. If they were not, then they would have discovered already that closed loop systems control and that this means that little of significance can be learned about these systems by studying input-output relationships (the traditional approach in psychology).

In 1976, just as I was starting to try to understand PCT, I ran into a book by Neisser (called Cognition and Reality) where he talked all about the fact that there is this "perceptual cycle" -- perception leads to cognition which leads to action which leads to new perception. It sounded great (at the time) but, it turns out, it was just words. Without a quantitative anaysis Neisser had no idea what his loop really meant -- it meant that perceptions were controlled relative to fixed or variable internal references. I gave a talk recently were someone cited Neisser's book as evidence that psychologists already knew what I was talking about. It's at times like that when I get very pissed at the penchant among psychologists to treat a turn-of-phrase as understanding.

Trust me. Whatever psychologists call their models, they are SR models pure and simple and they are treated this way. When psychology (and the other sciences of life) really catch on to PCT, it will be obvious, not from their words (they might call control theory something else) but from their actions; they won't do research in the way it is currently done. The IV-DV approach will disappear and you will start seeing studies of controlled variables using some version of "The test". Statistics will become unnecessary (unless the controlled variable is itself statistical) and experiments will be repeated regularly -- because they will always work (just like physics experiments). When you see that, then you will know that psychology has become the science of living control systems (as it should be). I don't imagine that we'll see that for some quite some time, unfortunately. But we'll hear a lot of people making up words that sound like control theory ("situated behavior" seems to be the latest contender). Remember, don't listen to what they say; watch what they do. When they start testing for controlled variables, then you know they've got it.

Best regards   Rick

Date: Sun, 9 Aug 1992 15:41:13 EST
Subject: disturbances in video game players

Although the video-game playing programs don't seem to have to cope with much in the way of disturbances, surely some could be added without too much drama. E.g., the Amazon dungeon might have its floor populated with conveyer-belts going in various directions & moving at variable speeds, and also, say rotating platforms (to make things more challenging, these also might be invisible to Sonja). The issue is then how much revision Sonja would require to cope with this. My suspicion is not much. E.g., if she's trying to go toward an amulet, the effect of a conveyer belt will be the same as if the amulet were moving, and the mechanisms already there ought to be able to cope with it (always move in the direction `toward the amulet', whatever that happens to be at the moment).

And then there's the issue of to what extent the continuously-varying-disturbance regime that exists at lower levels (wherein muscles are continuously varying the respose they emit to a given neural current, etc.) persists thru to higher ones. E.g. could you get a reasonable robot by hitching continuous control systems for the lower levels to a Agre-Chapman style one for the `tactics' of dealing with situations. The sheer fact that the existence of control is easy to miss is maybe evidence that this is really true.

Something else I'd like to find out is whether the `arbitration network' architecture used in the Pengi's and Sonja's central systems could be profitably redone along PCT lines, controlling for projected occurrence of desireable things and non-occurrence of undesireable ones. Two outcomes that would justify the effort, if they ocurred, would be:

a) more insight into the organization (knowing that such-and-such a actually a control system controlling for some value of some ought to help in relating these system's structure to their behavior

b) workable methods for the systems to improve their performance (there currently don't seem to be any).

An objection of Bill's that I don't buy is that the smarts in these systems come from the programmer: the main idea that the interative AI people had to shove off the deck was that you couldn't accomplish things in the world without elaborate advanced planning. It is therefore quite sensible to first produce systems that can function with zero or minimal planning, and then worry about how these architectures might arise in nature. This is basically Chomsky's point in generative grammar that it is futile to talk about the mechanisms of language-acquisition if you don't have a reasonably viable notion of what the outcome of acquisition is like.

Avery.Andrews

Date: Sun, 9 Aug 1992 08:14:35 -0600
Subject: models & demos

[From Bill Powers (920808.2000)]

My system has temporarily lost Kermit, so this will be sent late.

penni sibun 920807 --

>i've seen the crowd demo (thanks to avery) and i've seen sonja (pengi's
>daughter). quite frankly, you know, they're both just demos, and one >can

look at them, or even play with them, and assume they're video >games that aren't trying to demonstrate any particular theory of >activity. if you read the accompanying text and/or the code, you may >or may not speculate about whether or how the demos support the models.

Yep, that's the problem with just looking at behavior.

The crowd model could be set up with a single active person seeking a goal across the screen, and one stationary person in the way. When the moving person got close to the stationary one, the path would curve around the stationary person and head on to the goal. There are lots of ways to accomplish this result.

One way would be to write a program that says "Draw a line from the start position to within 20 pixels of the obstacle position, then decribe a semicircle centered on the obstacle, then continue the line to the position of the goal." That would be pretty trivial.

Another way would be to say "Move toward the goal position. If an obstacle appears within 20 pixels of the moving object's position, turn the path left and keep the obstacle 20 pixels away until you are again on a line between the starting position and the goal. Then move toward the goal position again until it's reached. Then stop." This is slightly better, but you still have to inform the model of its own position, the obstacle position, and the goal position by providing this data from outside the program. You also have to tell the model what "toward" means. You also have to tell it to stop when its distance from the goal is zero. You have to program into it the meaning of "turn left" and you have to provide a computation that will keep it 20 pixels from the obstacle. The BEHAVIOR of this model would be just like that of the first one, but it's not the behavior that we're interested in. It's how that behavior is created by the combination of the programmer and the running model. There's a little more going on inside this second model than inside the first one.

My way of programming an individual in the crowd program was done in still a different way. I told it which object was the goal and which were obstacles. I told it to use an avoidance control system for obstacles and a seeking control system for the goal. The avoidance control system had two sensors, one detecting the sum of all proximities to obstacles on the left, the other the sum of all proximities of obstacles to the right (weighted to reduce the effect of objects toward the rear). Proximity was defined as k/r^2, where r was the radial distance to a given obstacle. This created a perception similar to that of visual area subtended by an object at a distance r.

The avoidance system used both the sum and difference of left and right proximities of all obstacles. The error signal from the control system caused the direction of movement to curve to the left or right, curving faster as total proximity increased and simultaneously slowing the velocity. The destination-seeking system used the difference in left and right proximities of the goal-object; its error signal also made the path curve left or right, velocity increasing with error.

As a result, a person would start out toward the goal, veer left or right when it got near the obstacle and describe a path around it, then go on toward the goal in a straight line from the point where its avoidance path was aimed at the goal again. When it got to the goal it would slow down and stop. If you looked only at the BEHAVIOR, this model would look much like the first and second ones mentioned above.

Date: Sun, 9 Aug 1992 09:32:58 -0600
Subject: Re: disturbances in video game players

[From Bill Powers (920809.0900)]

Avery Andrews (920808) --

>The issue is then how much revision Sonja would require to cope with
>this. My suspicion is not much. E.g., if she's trying to go toward an
>amulet, the effect of a conveyer belt will be the same as if the amulet
>were moving, and the mechanisms already there ought to be able to cope
>with it (always move in the direction `toward the amulet', whatever >that
happens to be at the moment).

I'm not familiar with Sonja -- could you (or penni) give a precis of what the program does? And if possible, how it does it? How, for example, how would Sonja know what direction is "toward the amulet?" Would this be computed relative to Sonja's present heading? How would Sonja change that heading in the appropriate direction?

>E.g. could you get a reasonable robot by hitching continuous control
>systems for the lower levels to a Agre-Chapman style one for the
>`tactics' of dealing with situations. The sheer fact that the existence
>of control is easy to miss is maybe evidence that this is really true.

The lower control systems can control only for lower-level variables. If you posit a lower-level system that can perceive a target object and orient the body toward it, and another that can perceive the distance of the object and make that distance approach a reference distance, the higher-level systems don't have to accomplish those things. You can assume a reference signal that picks the target and another that picks the intended distance from the target, leaving it up to the lower systems to orient toward it and move the body toward it. But doing so will alter the relations of the body to everything in the environment (for example, doors, corridors, other objects). Something has to be monitoring these relations if they matter to the logical situation. If the motions in space are part of the simulation, the logic has to wait until the commanded result is perceived as having been accomplished; there's no point in going beyond "move to the amulet" until perception of that situation has been accomplished. Otherwise the logic won't pick up the consequences of physical interactions that affect achievement of the logical goals.

If Agee and Chapman are actually building logical control systems, that's great.

>Something else I'd like to find out is whether the `arbitration >network' architecture used in the Pengi's and Sonja's central systems >could be profitably redone along PCT lines, controlling for projected >occurrence of desireable things and non-occurrence of undesireable

> ones.

Controlling for projected occurrances sounds like model-based control. That is, the behaving system contains a model of the properties of the environment. In the imagination mode, it alters its actions (on the model) until the perceived result is the one wanted. Then the switch is thrown and the same actions are sent as reference signals to lower systems that actually carry them out. The error signals in the lower systems indicate the degree to which (and the way in which) the lower systems failed to produce the expected perceptions, and so shows how the model must be modified or the goal must be changed.

>An objection of Bill's that I don't buy is that the smarts in these
>systems come from the programmer: the main idea that the interative AI
>people had to shove off the deck was that you couldn't accomplish >things in the world without elaborate advanced planning. It is >therefore quite sensible to first produce systems that can function >with zero or minimal planning, and then worry about how these >architectures might arise in nature.

I agree that these AIers are moving toward the CT point of view. They don't, however, seem much interested in learning about the CT point of view (I tried). But until I see exactly how these guys would accomplish something like "move toward the amulet" I would still suspect that THEY know how to move Sojna toward the amulet, but SONJA doesn't. Maybe this doesn't matter. But maybe it does.

To me, the critical question is "what can Sonja know and when can Sonja know it?" If that's not clear in the model it has to be made clear. Not all logical decisions can physically be carried out. None of them is carried out instantly, and none can be carried out without altering the situation on the way to accomplishing them. If the logic simply runs blithely along assuming that every commanded result is immediately and successfully brought about, or if the model itself isn't detecting the current state of the environment in the relevant regards, this isn't a CT model. But tell me more about it.

Best   Bill P.

Date: Sun Aug 09, 1992 3:42 pm PST
Subject: Sonja

[Avery Andrews 920810]
 (Bill Powers (920809.0900))

At the moment I don't have time to write up a thorough rundown of Sonja (insofar
as I understand her), but Chapman has written a book:

  Vision, Instruction and Action, MIT Press (1991 (I think)). c. $35.

What she does is play a video game (Amazon, based on a commercial game called `Gauntlet'. Something I hadn't registered when I wrote my earlier postings was that being a video-game player, Sonja only gets to move the Amazon in the usual eight joystick directions, but I don't think this is an essential limitation of the approach.

As I recall, if you tell her to `get the amulet' (she not only plays the game, but takes advice on how to do it, since she's not very smart), she finds an amulet on the `screen' (like a human player, she has a top-down view of the world she's moving the amazon icon around in), runs a `ray' from her present position to it, & follows the ray to the amulet (with kludgy patches for obstacle-avoidance. Interestingly, something she's rather bad at is navigation: I bet her navigational abilities could be substantially improved by borrowing the Crowd code to send out a fleet of `pathfinders' seeking the amulet, & following the course of the one who gets there first.

Her visual system is a very serious attempt to be neurologicallay realistic (based on Simon Ullman's visual routines theory), although the fact that it had to run on serial hardware meant that there had to be a certain amount of cheating.

More later, perhaps ... non-PCT duties are beginning to catch up with me. Plus I don't want to do too much more raving about these `reactive planning' systems until I've put together a (even very simple) one.

Avery.Andrews@anu.edu.au

Date: Mon Aug 10, 1992 6:51 am PST
Subject: thanks

from: eric harnden(920810)

i want to waste a litle bandwidth and thank the group, and in particular mr nevin and mr powers, for providing me with such stimulating and enlightening discourse over the past week. i have always enjoyed the conversational level of this group, and have benn especially rewarded recently with a wealth of material that illuminates my worldview (while at the same time providing the impetus to remain in the darkness of my study, to read and program). i'll be on vacation for the next two weeks, and will be taking the printouts of this last week's discussion with me. again, thanks. you've given me a lot to think about.

Eric Harnden

Date: Tue Aug 11, 1992 8:06 am PST
From: Dag Forssell / MCI ID: 474-2580

TO: Dick Robertson (Ems)
EMS: INTERNET / MCI ID: 376-5414
MBX: urrobert@UXA.ECN.BGU.EDU
Subject: Plooij
Message-Id: 75920811160657/0004742580NA2EM

Visited UCLA library yesterday. The Journal of Child Psychology & Psychiatry does not yet have the article. (May & July). The other journal is under consideration for subscription.

Can you please let me have copies of the two articles?

Have made progress on my project on behavior of perception, by laying out the graphics part. Now working (in between talking to companies) on the descriptions.

Trust you had a good flight. We spent Monday with Brent Dennis visiting Mesa Verde.

Regards,

Dag Forssell
23903 Via Flamenco
Valencia, Ca 91355-2808
Phone (805) 254-1195 Fax (805) 254-7956
Internet: 0004742580@MCIMAIL.COM

Date: Tue, 11 Aug 1992 12:10:22 +0200
Subject: Re: Models and demos -- what sells

From Oded Maler 920811]

[Rick Marken (920808) about SR <--> Control]

I think I start to understand the reason for some ongoing misunderstanding. You probably agree that some parts of the control loop when, seen in isolation are S-R, that is, some input goes in and some output goes out which depends on the input. And there is a mechanism responsible for it, it takes reference and perception, say, and produces force. This is *all* what roboticists mean by S-R. So many of their systems as you admit are control systems (maybe of a simple kind) although they might use terminology such as S-R, reinforcement etc.

On the other hand, in psychology, these terms are much more loaded, and in particular, they come with a "methodology" based on statitistical tests etc., which are supposed to reveal the underlying mechanisms from the outside without opening the black box. Your insight on control suggests that this is useless, and probably, as someone educated in psychology, this was one of the great revelations of your life. But for engineers (unlike reverse-engineers) this problem does not exist, because they build the boxes, they know the I-O ("S-R") characteristics of (some of) the components they put in the loop, and all your protests against S-R, which might be valid in the context of psychological experiments, seem like fighting non-existent windmills. They don't know and don't care about ID-IV tests and the only arguments you may have with them is whether it is useful or not to use the language of hierarchical servo loops, reference signals etc. in order to achieve interesting performance (I refrain from using the word "behavior"..).

Best regards --Oded

Date: Tue, 11 Aug 1992 13:01:43 +0200
Subject: On worlds - Blocks and real

[From Oded Maler 920810]

{Since I received no feed-back (from the list-server..) on this post from yesterday, I re-send it. I apologize if eventually it will be posted twice}

(Concerning the recent exchange on simplifying assumtion, disturbance-free worlds, modeling etc.)

[In this short unfocused essay I use often the term "Blocks world" which is a common metaphor characterizing symbolic AI toy worlds, where the world consists of objects,(block-a, block-b), properties (green(block-a)), relations (on(block-a,block-b)) and operators that transforms the state of the systems (e.g., remove one block from another). This gives the flavor of the world on which classical AI problem-solving took place.]

****

Suppose for the moment that the object of study of AI, AL, etc. is not defined with any reference to any reality of living systems. It is the study of mathematical platonic entities obeying some rules, interacting with some environment which is also an abstract object with its own formal dynamics. I'm sure that at this level no one will object the validity of various results obtained there (only their relevance to something interesting). I'll give two concrete examples one in AI flavor and one in AL:

1) An algorithm, that given a "goal" (in this setting, a representation of some configuration of the "environment") produces a sequence of "actions" that leads from any state to this state.

2) A set of simple local dynamical rules, that when employed by many similar abstract entities leads to an interesting emergent global "behavior" (=trajectories in the state-space).

Now, we try to see whether there is any relevance to the "real" thing. Considering the first example, we all know (by now) that the world does not come with labels "a" and "b", and making sense of the sensory input is a big mess. We also know that any abstract action that we want to perform should be realized in a smaller scale with disturbances etc. In other words, the real world is not the Blocks-world. On the other hand, we feel from introspection that *some* of our mental activity is done in this (graph-searching, if-then-else) mode, and *if* the lower-level systems take responsibility for all the messy details, uncertainties, and disturbance, it can transduce the real-world into a Blocks-world, and then a higher-level system can operate at this abstract level, and thus the achievements of systems of type (1) above can become relevant.

This was (and probably is) the tacit assumption of symbolic AI, cognitive psychology, etc. and without it all this research is rather irrelevant to human mentality and of little relevance to engineering (except for some non-robotics AI such as [non-video] game playing).

There might be two objections to this view:

I) It is impossible to construct such a real-to-abstract transducer without solving the unsolvable categorization problem (and certainly not in real-time).

II) - And this is the most intriguing fact for me - *EVEN* if such a translation exists, except for some small toy problems, the planning problem in such a world is computationally intractable.

All this means that the blocks world is uncomparable in its complexity to the real world (as perceived by our problem-solving level): in one sense it is simpler because its availiability to us subsumes a non-existent unrealistic transducer, and on the other hand it is more complex because many problems that we can solve with our slow brains, can not (provably) be solved in the blocks world using the fastest computers.

The second observation might suggest that concerning negative results, simplifying assumption are useful bacause it demonstrates that *even* in a disturbance-free environment, something cannot work - so it certainly won't work when ideal transducers are replaced by noisy ones.

I've been attracted to the bottom-up approach (e.g., Powers, Brooks) hoping that some principles discovered while going up in the levels will explain how, after all, it is possible to plan and solve problems in the real-world in the higher level. In what way does the fact that the symbols on which the higher-level systems work, are "grounded" (to use a popular term) in lower-level interaction with our particular concrete world, make problem-solving easier than is some arbitrary abstract world (and this is also, I think, related to some of the ecological ideas of Gibson).

Ok, I see I'll not touch the AL example this time. Feed-back (of any polarity) is wellcome.

--Oded

Date: Tue, 11 Aug 1992 08:14:07 EDT
Subject: Thar she blows! PR sounding

[From: Bruce Nevin (Tue 920811 08:03:14)]

There is potential here for some broader exposure. An angle that would do very well with this audience is the "staid conservatives ignore revolutionary alternative because it contradicts their expectations" story. Something like this (but with less basis) played very well in Scientific American recently regarding language history and linguistic reconstruction.

Thar she blows! Shall we attempt to take him, Cap'n?
I don't know, mate. We're only rigged for cod . . .
But think of what he'd fetch if we succeed!
Yep. Might stove us in too.

> X-To: Cybernetics + Systems list
<cybsys-l@bingvmb.cc.binghamton.edu>

> Really-From: Carl Zimmer <0005399004%mcimail.com@BINGVMB.cc.binghamton.edu>
> Date: Fri, 7 Aug 92 21:07 GMT

> I just signed onto this list and wanted to introduce myself. I'm associate
> editor at Discover magazine. The subjects covered here are exactly the
things
> we love covering, so I'll be glad to hear from people. I may want to send a
> note in the future to find people involved in a particular cybernetics
branch,
> but I'll try not to be too obnoxious. If anyone wants to reach me, they can
> get me directly at MCI Mail (ID: 539-9004).

> Carl Zimmer

Bruce Nevin bn@bbn.com

Date: Tue, 11 Aug 1992 09:19:21 -0600
Subject: Boocks worlds

[From Bill Powers (920811.0900)]

Oded Maler (920811) --

I like what you're saying about the difference between the engineering problem with PCT and the traditional psychological one. I don't have any basic objection to the blocks-world approach -- we play that game ourselves, in a way. But I do think that the blocks world can lose its bearings when it isn't tied to the realities of real living systems as closely as feasible.

Ask yourself this question (I don't know the answer and would be interested in hearing it). What would happen to AI and AL models if all computations were limited to 1 percent accuracy? I know that Ashby got into trouble with the implied problem when he abandoned control theory and went in for the compensation model. He ended up using examples in which all the variables took on small integer values, so it was possible to subtract 7 units of output of the compensating system from 7 units of disturbance and get 0 effect on the critical variable. This infinite precision conceals a lot of problems with small cumulative disturbances, nonlinearities, and the like. In a similar way, the "motor program" people, who assume open-loop calculations of output signals, use solutions of simultaneous differential equations in which multiple time-integrations have to be calculated with full available floating-point precision in order to work even approximately -- if these models were required to behave in a variable environment (no matter how predictable) for a couple of hours or days without being reset, they would surely crash.

Even at the logic levels, it seems to me, the AI and AL types of models assume too much precision: the models never make a logical error as human beings do. This, it seems to me, shows that the kind of logic they're using, which is computer logic or Boolean algebra, is not the kind that the real system uses. If you build an organization that can work only if no computational errors ever occur, you never come up against the basic problem of error-detection and error-correction, and so never realize that it MUST be solved.

The only solution I know of to the precision problem is to use closed loops, reference signals, and so on. Even fuzzy-logic models, which don't control very well, use closed loops. It isn't only the categorization problem that's intractable in an open-loop model. It's the control problem at every level. What, other than sudden death, tells the system that its actions aren't producing the calculated results, or that its perceptions aren't representing the world in a useful way? What tells it that while its logic is all right, its premises are screwed up?

I don't think that the digital computer revolution has done behavioral science any favor. The digital computer was designed to work with absolute reliability, so that it always does what you tell it to do. Even the stepper motor was invented to make S-R behavior feasible. If you can count steps, and if the circuitry assures that a step is never gained or lost, and if the environment never changes unpredictably, you don't need any feedback. I think that lots of models of behavior assume this absolute precision of action that is inherent in a digital computer -- that's why they make so little use of feedback, why they seem to take action so much for granted, and why they ignore the subject of disturbances. Their inventors work in a world where if you say "add two and two" you always get 4.0000000000000000000.

Best Bill P.

Date: Tue, 11 Aug 1992 16:16:26 -0500
Subject: MCI Mail Link Down

[from Gary Cziko 920811.1615]

Just to let you know that since around Sunday the MCI Mail gateway with Internet has not been working. This means that no posts to CSGnet have been getting to:

00004742580@MCIMAIL.COM FORSSELL Dag, Valencia, CA
00004912499@MCIMAIL.COM CLARK Bob, Cincinnati, OH
00004972767@MCIMAIL.COM WILLIAMS Greg, Gravel Switch, KY
00005096370@MCIMAIL.COM FOLSON Hank, Henry James Bicycles

As I understand it, nothing from Internet or Bitnet can get to any MCI Mail subscribers now, although they can apparently send messages out.

Date: Wed, 12 Aug 1992 09:47:47 EST
Subject: PCT and Sonja architecture

More on Sonja. What I'll talk about is killing monsters. Sonja's architecture is a 3-level scheme of `registrars', who notice interesting features (`aspects') of what's happening, `proposers', who decide that it would be nice to do something, and `arbiters', who settle conflicts between proposers. So to make it worthwhile to kill a monster, a registrar has to decide that it is dangerous. This is achieved by having an autonomous process put a `marker' on the closest marker (a marker can be thought of as a register connected to the visual system, which automatically tracks something, and from which a certain amount of information about that something can be extracted). When the marker returns the information that the monster is within a certain distance, the `the-monster-is-dangerous' registrar gets activated (think of the registrar as a perceptual circuit).

What happens next depends on peculiarities of video-games. The Amazon icon (the little figure that Sonja maneuvers thru the dungeon) can only move and shoot in the usual eight joystick directions. So a natural way to represent the full range of directions is as pair of vectors, chosen from these eight, the larger one called the `major component', the other the minor. What you need to do to shoot something is to move till the minor component is zero (`align' with the monster), then face in the major component, then fire. So there is a proposer `align-with-the-monster', which gets activated when `the-monster-is-dangerous' registrar is, and which proposes moving in the direction of the minor component of the direction to the monster. The effect of this is of course to reduce the minor component, so `align-with-the-monster' is basically the comparator of a control system controlling for the following conditional:

```
   if
      the-monster-is-dangerous
   then
      the-minor-component-of-the-direction-to-the-monster-is-zero
```

Note that since the system works by closed-loop interation with its environment, if the monster were to cease to be dangerous, (by virtue of being killed by someone else, or of running away), Sonja would stop trying to align with it. Also if it moves, the system will track it. Also that it doesn't matter if the representation of distance is low-precision, as long as its pretty good near zero (so you know when your actually aligned) (I imagine a function on the interval [0,1) would do just fine, asymptotically approaching the upper bound).

I think this kind of conditional control system is sort of interesting, and the way I'd suggest building them is like this:

```
                               | Ref. sign.
                               |
      --------------------------------------|
      |                        |            |
      |                        v            v
      |                     --------     ----------
      |          ---------->| Comp |---->| *-Gate |
      |          |          --------     ----------
```

```
   --------   -------                      |
   | Cond |   | Perc |                     |  Error (influencing proposed
   --------   -------                      |    action)
    ^  ^        ^  ^                       v
    |  |        |  |
    |  |        |  |

    Lower level Perceptions
```

`Perc' is the perception that is to be held at a scalar value, given that condition detected by Cond olds is acheived, & this is accomplished by means of the multiplicative gate not passing the comparator's error signal unless the input from Cond is nonzero. The Cond boxes in fact look to me as if they ought to be able to do the job of the arbiters, tho maybe I'm missing something here.

When the alignment is good enough, the another proposer, `face-the-monster' will kick in, which can be thought of as a control-system that tries to zero the difference between the-direction-the-amazon-is-facing and the-direction-toward-the-monster. Given the way the Amazon game works, there aren't any disturbances for this system to cope with, but it would be easy to make some up (e.g. whirlwinds that would tend to spin Sonja around). But face-the-monster should only fire up when `the-minor-component-of-the-direction-toward-the-monster-is-zero' fires up, so, in effect, its Cond box tests for the condition that the error signal of `align-with-monster' is (almost) zero. Finally, when `face-the-monster's error signal is zero, `shoot-at-the-monster' is fired up, which might be thought of as a system whose reference level is to perceive the monster to be terminated (Sonja shoots until the monster disappears).

So the whole architecture looks to me to be full of control systems, and rather like my `getting a beer' scenario (as Penni pointed out to me some months ago). But it is *not* presented this way by Chapman (nor is Pengi so-presented by Agre), so one issue is whether the PCT approach is really a more helpful way of thinking about these systems than whatever their authors are doing (I find their actual discussions of the control architecture to not be especially helpful, and the PCT perspective to be useful in figuring out what's going on).

A final thought: Chapman himself is not interested in Sonja's control architecture, on the basis that deciding what to do is easy: what's hard is perception. So he would probably not be interested in any improvements that the PCT perspective might offer.
==
Date: Tue, 11 Aug 1992 22:07:08 -0600
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: "William T. Powers" <POWERS_W%FLC@VAXF.COLORADO.EDU>
Subject: Sonja; dues

[From Bill Powers (920811.2000)]

Avery Andrews (920811) --

That is a most informative post about Sonja. Your attempt to cast what's going on in PCT terms makes it so obvious that there are control systems there that I wonder why Chapman doesn't see them. On the other hand, I think Chapman is managing to get some orderliness into his scheme, which is to be admired.

Your proposition about the gated control system is a good start. Remember, though, that in a neural model, to turn off a one-way control system (no negative error signals) it is sufficient to set the perceptual signal to zero (negative feedback alone can't cause an error signal if the reference signal is zero). The implication is that the sensing of the "condition" (a logical condition) is the input function of a higher-level control system, the output of which sets the reference signal for the "scalar control system." The higher system, being program-level, can actually select which lower-order perceptions to adjust by choosing which reference signals to bring above zero. Rick Marken has a logic level in his spreadsheet demo program.

>... to make it worthwhile to kill a monster, a registrar has to decide
>that it is dangerous. This is achieved by having an autonomous process
>put a `marker' on the closest monster (a marker can be thought of as a
>register connected to the visual system, which automatically tracks
>something, and from which a certain amount of information about that
>something can be extracted). When the marker returns the information >that
the monster is within a certain distance, the `the-monster-is->dangerous'
registrar gets activated (think of the registrar as a >perceptual circuit).

This is verging on hierarchical perceptions. A "marker," for example, clearly has to include a perceptual function if it's to "return information that the monster is within a certain distance." The recognition of the monster is one perception; the distance is another. Those perceptions becomes an input to a registrar. Why not just say that there's a system that senses and controls the distance of a particular object? This would provide for approaching or avoiding any object of whatever kind by adjusting the reference-distance signal. If a higher system classifies this object as dangerous, it can set the reference signal for proximity to a low value to avoid it, or to a high value if the decision is to attack it. If the object is classified as desirable, the same reference signal can be adjusted formaximumn proximity in order to approach the object (so as to do whatever one does with desirable objects).

>When the alignment is good enough, the another proposer, `face-the-
>monster' will kick in, which can be thought of as a control-system that
>tries to zero the difference between the-direction-the-amazon-is-facing
>and the-direction-toward-the-monster.

If the amazon is given real perceptual functions, the-direction-the-amazon-is-facing won't be perceived. The amazon will simply perceive everything else relative to the direction of looking. To align with something, the amazon simply turns and translates until the object to be faced is centered. The calculation of what the amazon can see is part of the model of the environment, not of the behaving system.

The problem here is that it's Sonja, not the amazon, that contains the control systems. So Sonja is acting like the game player, being above the fray looking down on it, using her own properties but trying to make the amazon do things. This is analogous to Chapman sitting up in the sky watching and describing and evaluating what Sonja is doing. Sonja's properties get all mixed up with the properties of the amazon. The means by which the amazon kills a monster are not the means by which Sonja makes the amazon do things that Sonja interprets as killing the monster (for Sonja, making it disappear). Sonja's proposer that says "face the monster" does not make SONJA face the monster, because Sonja is already looking down on the monster and can see it. This proposer is really proposing that an icon on the screen be rotated and translated until a line projected along its heading passes through the monster and corresponds to a joystick direction (two conditions to be controlled).

I think that Chapman has a point-of-view problem here. He is actually trying to solve a very much more complex problem than he has to solve. He's trying to represent not only the game problem, but the player's identification with the amazon, and the stories that he himself tells about Sonja as she goes through the motions, and about the amazon as the action proceeds. Those stories are Chapman's, not Sonja's. Unless he wants to give Sonja the ability to make up such stories, they aren't part of the Sonja model. If Chapman wants to model the game player, then he should model what the game player is actually controlling, rather than jumping around among Sonja's point of view and the amazon's and his own. The verbal embellishments only make it harder to see what is actually being controlled. The monsters on the screen are not "dangerous." They are simply icons in certain spatial relationships with the amazon. Those spatial relationships can be controlled relative to any reference-relationships Sonja pleases. The task of a higher-level system is to prevent certain relationships from occurring and to achieve other relationships. In order to do either, it's necessary that the relationships be under control. Once they're under control, particular sets of them can be named, and a higher-order system can worry about maintaining certain sequences of named relationships and certain logical conditions in which the named relationships are elements.

Am I getting this across? I hope so, because this is as close as I've ever come to saying what I find wrong with this sort of modeling. This is what I mean when I say that too much of the modeler is outside the model making it work. There's nothing in the model of Sonja that could come up with the label "dangerous," with all its meanings. The mechanism for doing that is inside Chapman, as are the meanings. The best Sonja can do is avoid a logical condition involving a configuration and a relationship: (monster AND monster-close-to-amazon). From Chapman's point of view, if Sonja does indeed manage to keep that condition false, it will seem as though Sonja perceives the monster as dangerous and acts appropriately. But that's only a description; it's not what makes Sonja work.

In the crowd program, an individual will sometimes go through an opening into a pocket of closely-spaced obstacles. The moving person will turn this way and that, going through all sorts of loops, and eventually leave through the same opening. It looks as if the person is perceiving this-is-a-way-through-the-obstacles, then deciding there's-no-way-out-of-this-pocket and therefore i'd-better-look-for-the-way-I-came-in and finally doing go-the-other-way-through-the-same-opening. If we set this up as a logical problem with markers and proposers and arbiters, we could probably reproduce something like this behavior -- but in fact nothing like that is actuually happening, or is actually necessary. Setting up the problem in terms of those hyphenated phrases leads into doing something quite simple in a horribly complex and unlikely way. That's because the phrases bring in a whole background of the OBSERVER'S thoughts and associations and descriptions, none of which the model is capable of perceiving or producing. What you end up with is an extremely complex model that produces simple behavior. What we want, of course, is a simple model that produces complex behavior. I don't think that the approach taken by Chapman and Agre is going to create a model simpler than the behavior it explains.

But all this is still based on second-hand information. I've applied through interlibrary loan for Chapman's book. When it gets here I'll have another look and see if I still come up with the same critique.

Bruce Nevin and others who are students: At the CSG meeting, it was proposed to increase student dues to $10 per year to cover the costs of Closed Loop. A counterproposal was that we increase ordinary membership cost from $40 to $45 to subsidize the students' subscriptions. The counterproposal carried. Those who have already paid their full $40 for 1992-93

Best Bill P.

Date: Tue, 11 Aug 1992 21:33:10 PDT
Subject: Re: Sonja; dues
In-Reply-To: "William T. Powers"'s message of Tue, 11 Aug 1992 21:07:08 -0700 <92Aug11.210927pdt.11917@alpha.xerox.com>

(penni sibun 920811)

i've got lotsa comment on this thread, but i've been busy and my typing capacity is limited. however, i saw an opportunity to make a simple answer.

> [From Bill Powers (920811.2000)]
>
> The problem here is that it's Sonja, not the amazon, that contains the control systems.

this may be a problem and chapman may have problems, but this isn't one of his. chapman's programs, pengi and sonja, are programs that play video games. it does get occasionally confusing, but that in fact is the nature of video games: the player identifies w/ the character in the game. (that's why *we* get confused: i don't know if sonja confuses itself with the amazon or pengi confuses itself with the penguin.)

chapman (1991:59-61) explains the rational for video games clearly:

pengi was...designed to demonstrate concrete-situated ideas about activity. these ideas emphasize perceptual connection with the world,

so i wanted a domain which required substantial perceptual
interaction, but which would not require that i spend years solving
known difficult probleems in vision research. video games have just
these properties....because video games are two-dimensional and
because the images involved are generated by a computer, i could bypass
hard problems in early vision such as noise tolerance, stereo, and
shape-from.''

you may or may not consider this a reasonable choice, but chapman made
it knowledgably.

   Am I getting this across? I hope so, because this is as close as I've ever
   come to saying what I find wrong with this sort of modeling. This is what I

can you try again, with the above confusion cleared up?

--penni

Date: Wed, 12 Aug 1992 14:55:14 +0200
Subject: Re: Blocks worlds

[From Oded Maler 920812]

(Bill Powers, yesterday):

> What would happen to AI and AL models if all computations
> were limited to 1 percent accuracy? ... etc.

The point is that there is some level where we (seem, at least, to)
work in discrete terms, e.g., "if my PC doesn't boot then I replace
the diskette" ,etc. The underlying assumption is that there are
lower-level modules, employing feed-back all the way down, making
sure that the abstract action "replace the diskette" is being performed
in a noisy world (in fact this is a too complicated example because
even sub-tasks of this action are not yet explainable *in detail* in
PCT style - try running Little Man in a non-uniform environment with
obstacles, non-reversible consequences, dead-ends [and this is by
no means an attempt to dismiss its achievements - just to note that
there are many hard problems in the highr-levels]).

Now, it's not clear what 1% accuracy means in this abstract world,
because the disk-drive cannot be 99% pregnant - and if I buy from you
a low-level black box, I trust that it will work. If you mean
inaccuracy in the higher-level, that is, replacing the diskette cannot
be performed or it may not solve the booting problem, and maybe
kicking the table will solve it, then you come to the realm of the AI
methods - searching among alternatives for actions, an activity that
needs a lot of creative imagination (more than I have for the moment)
in order to be perceived as gradient following, as servoing for
complex perceptions etc. The people who investigate such phenomena
surely abstract away the lower levels in the same way you abstract
away the micro-structure of neural and muscular cells.

Now I come back to my basic rhetoric question: goal-achieving in this abstract world with simplifying assumptions should be much easier than in reality - the lower-level black boxes are assumed to be perfect, there is no conflict in actions, and still it is very hard to to give tractable way of solving everyday problems. Why? What is wrong in our verbal account of our introspection concerning the way we operate in this level? Do the higher-level percepts and references carry with them some concrete features of the lower level ones from which they are constructed/grounded in such a way that the solution "flow" instead of being searched-for/calculated?

I don't expect immediate answers..

Best regards --Oded

Date: Wed, 12 Aug 1992 09:15:28 -0600
Subject: AI & AL problems

[From Bill Powers (920812.0800)]

penni sibum (920811) --

I don't fault Chapman for bypassing problems with vision, form recognition, noise tolerance, and so on. Even if we knew how a game player converts an array of pixels into a signal standing for "monster," we might prefer to leave out those computations just to fit the simulation into a computer.

>...it does get occasionally confusing, but that in fact is the nature >of video games: the player identifies w/ the character in the game. >(that's why *we* get confused: i don't know if sonja confuses itself >with the amazon or pengi confuses itself with the penguin.)

What I'm trying to get at is that the *model* is not confused; it does what it does. When you boil down all those hyphenated phrases, they have to become variables and operations on them. All the verbal interepretations are going on outside the model, in the modeler. I'll have more to say after I see Chapman's book.

Oded Maler (920812) --

>The point is that there is some level where we (seem, at least, to)
>work in discrete terms, e.g., "if my PC doesn't boot then I replace
>the diskette" ,etc.

I don't think that the discrete terms are always as discrete as they seem. Suppose you do stick a disk in drive A and turn on the computer. You hear the fan start, the red light on the disk drive blinks, the disk drive whirrs and buzzes ....

At this point, has the PC booted? What is the state of the categorical/sequential/logical perception, "booting up the pc?"

It goes on. The B drive light blinks and the drive buzzes, the hard disk light flickers, the printer goes kachunk, and finally all the action quiets down. Only one thing -- during all of this, there's been nothing on the screen.

Now what's the status of "the disk drive is booting up?" All the time that the familiar activity was going on, you were getting perceptual signs that the computer was on its way up (a familiar event). But with nothing on the screen, part of the process is missing. The computer seems to be booting up, but it's not doing it quite right. When the action stops, you've got a big question mark hanging: did it really boot up or not? What's wrong? You perceive MOST of "the computer is booting up" but not ALL of it. Eventually you realize that there's no pilot light on the monitor, and you turn the switch to "on" -- and the screen fills up with the usual gobbledygook. NOW the computer has booted up.

I'll bet that Martin Taylor would say that the perception of the discrete category or state is probabilistic. When something is missing, the probability-value of the perception is less than 1. I would just say that the perception changes on a continuous scale from none to the reference amount. There are probably other ways to interpret this experience, implying still other propositions about how the perception of discrete categories works. But the one way of putting it that we wouldn't use, I think, is "either the computer is booting up or it's not." That strictly mathematical interpretation doesn't fit the way perceptions actually behave. When everything happens right except for the screen display, we don't experience this as "the computer is not booting up." We experience it as "something's wrong with the way the computer is booting up."

To me, this means that set theory, symbolic logic, and computer programming don't quite fit the way the brain works at these levels. No doubt there are aspects of what the brain is doing that resemble the way mathematical representations behave, but I think it would be a mistake to assume that the mathematical processes are anything more than an approximation to the real system's mode of action. It's possible that something else entirely is going on.

Think of the difference between "absolutely" in the following sentences:

"What you say is absolutely true"

"I'm pretty sure that what you say is absolutely true."

The meaning of "absolutely" is weakened by saying "I'm pretty sure." It's not negated or affirmed; it's left somewhat less than affirmed. While the words still seem categorical, their meanings are shifted in a continuum, as are the meanings of anything further we say on the subject:

"I'm pretty sure that what you say is absolutely true, but let's act cautiously."

Computers can't act cautiously; all they can do is act or not act.

>... try running Little Man in a non-uniform environment with obstacles,
>non-reversible consequences, dead-ends [and this is by no means an
>attempt to dismiss its achievements - just to note that there are many
>hard problems in the highr-levels]).

Oh, yes, there certainly are. I think, though, that as we build up a working model level by level, the nature of those hard problems will begin to look different. Maybe the hard problems will prove to have simpler solutions. Look how hard reaching out and pointing seems if you try to accomplish it without feedback, the way the motor program people are doing. In the motor program scheme, the higher levels have the burden of analyzing the environment and the dynamics of the arm, and computing signals that will result in a particular trajectory and end-point. It would take at least a program-level system to compute the inverse dynamics of the system. With the control-system approach, the higher levels don't have to be concerned with that at all: all they have to do is decide on the desired relationship between the fingertip and the target. So all that terribly complicated computation disappears from the model. You're still left with problems, such as what to do when there's an immovable obstacle in the way of pointing, but I'd rather try to solve problems like that than the problem of computing inverse dynamical equations.

>Now, it's not clear what 1% accuracy means in this abstract world,
>because the disk-drive cannot be 99% pregnant - and if I buy from you
>a low-level black box, I trust that it will work.

The problem of 1% accuracy shows up when a logical process commands a result to occur, and it occurs with a 1% error. This is a reasonable error for a good lower-level control system that's dealing with moderate disturbances. What is the next process in line to do? If it just proceeds, it will begin with a 1% error, and after it's done, the error might be 2%. There's nothing to prevent a random walk as these errors accumulate, until what's actually accomplished in response to a command or a decision bears no resemblance to what was intended at the higher level. Somehow these little quantitative errors have to be taken into account at the higher, supposedly discrete, levels. Of course in what you term an abstract world, such discrepancies can't exist; there's no provision for them. This may mean that the abstract world is not a good representation of the higher levels of perception and control. Of course human beings can behave as if the world were discrete -- but the world they deal with then is imaginary, and when the abstract processes are required to work in the real-time world, they will probably fail.

I'm probably trying to make my case too strong. A lot of behaviors governed by discrete decisions aren't critical; if I decide to drive into Durango, it doesn't much matter whether I hit the center of town within 1%. If I miss the Durango Diner by one block I can just park and walk the remainder of the distance (if I can find a spot that close). The lower level systems can make up for a lot of imprecision and idealization on the part of the higher ones -- they fill in the critical details that the "abstract" systems leave out.

One last observation. Control systems don't deal with pregnancy -- only with perception of pregancy. If a woman you know well shows up looking considerably stouter than usual, you could easily perceive her as 60% pregnant and 40% overweight. You won't congratulate her, and you won't form an opinion of her eating habits, because there are two conflicting perceptions of her state. To assert that a woman can't "be 99% pregnant" is to make an epistemological assumption, which is that the world is identical to our perceptions of it. We assume that the world can't be in two mutually-exclusive states at once; in fact, that's the essence of categorical thinking. But the world we experience can be in states between different categorical boundaries and between the logically true and false.

>What is wrong in our verbal account of our introspection concerning the >way we operate in this level? Do the higher-level percepts and >references carry with them some concrete features of the lower level >ones from which they are constructed/grounded in such a way that the >solution "flow" instead of being searched-for/calculated?

Higher systems are both continuous and discrete, I think. Anything a person does is an example of human systems at work. People can obviously get into a "calculation" mode where they apply rules literally and arrive at yes-no results. So a model of a brain has to contain the ability to do this. On the other hand, these processes can also operate in a sloppier way where the processes look more like a flow than a switch. So the model has to be able to do that, too. I think that at what I call the program level, there is a lot more going on than digital calculation. I think this is a generalized computer that can follow ANY kinds of rules -- the rules of Boolean logic are just one kind.

>I don't expect immediate answers..

Me neither.

Bruce Nevin (920811) --

Do you think we should send that guy at Discovery some copies of Closed Loop and invite him to join this net? Why don't you ask him?

Best to all, Bill P.

Date: Wed, 12 Aug 1992 10:59:26 -0500
Subject: Re: Thar she blows! PR sounding

[from Gary Cziko 920811.1816]

Bruce Nevin (Tue 920811 08:03:14) let us know that an associate editor for _Discover_ magazine, Carl Zimmer, has joined the CYBSYS-L group and is interested in cybernetics topics.

I would be happy to zap him with a bunch of CSG-PCT stuff. But I will have to wait until MCI Mail gets its Internet gateway working again. Any reservations out there?

> Thar she blows! Shall we attempt to take him, Cap'n?
> I don't know, mate. We're only rigged for cod . . .
> But think of what he'd fetch if we succeed!
> Yep. Might stove us in too.

And CSGnet, CSG, and PCT is NOT just rigged for cod (apologies to Bill Silvert in Nova Scotia)--I think we are able to and willing to take on the whole enchilada (as Rick Marken in LA would say).--Gary

Date: Tue Aug 11, 1992 1:18 pm PST
Subject: Your removal from the CSG-L list

Dear subscriber,

As of Tuesday, August the 11th of 1992, you have been removed from the CSG-L
distribution list (Control Systems Group Network (CSGnet)) by "CZIKO Gary, U of
Ill., Urbana" <CZIKO@UIUCVMD>.

Virtually, The LISTSERV management

Date: Wed Aug 12, 1992 8:32 am PST
Subject: Your subscription to list CSG-L

Dear networker,

As of Wednesday, August the 12th of 1992, you have been added to the LISTSERV
distribution list CSG-L (Control Systems Group Network (CSGnet)) by "CZIKO
Gary, U of Ill., Urbana" <CZIKO@UIUCVMD>.

You may leave the list at any time by sending a "SIGNOFF CSG-L" command to
LISTSERV@UIUCVMD. Please note that this command must NOT be sent to the list
address (CSG-L@UIUCVMD) but to the LISTSERV address (LISTSERV@UIUCVMD).

The amount of acknowledgement you wish to receive from this list upon
completion of a mailing operation can be changed by means of a "SET CSG-L
option" command, where "option" may be either "ACK" (mail acknowledgement),
"MSGACK" (interactive messages only) or "NOACK".

Contributions sent to this list are automatically archived. You can obtain a
list of the available archive files by sending an "INDEX CSG-L" command to
LISTSERV@UIUCVMD. These files can then be retrieved by means of a "GET CSG-L
filetype" command, or using the database search facilities of LISTSERV. Send an
"INFO DATABASE" command for more information on the latter.

Please note that it is presently possible for anybody to determine that you
are signed up to the list through the use of the "REVIEW" command, which

returns the network address and name of all the subscribers. If you do not wish
your name to be available to others in this fashion, just issue a "SET CSG-L
CONCEAL" command.

More information on LISTSERV commands can be found in the "General
Introduction guide", which you can retrieve by sending an "INFO GENINTRO"
command to LISTSERV@UIUCVMD.

Virtually, The LISTSERV management

Date: Wed Aug 12, 1992 8:51 am PST
From: g cziko
EMS: INTERNET / MCI ID: 376-5414
MBX: g-cziko@uiuc.edu

TO: Robert K. Clark / MCI ID: 491-2499
TO: * Dag Forssell / MCI ID: 474-2580
TO: Hortideas Publishing / MCI ID: 497-2767
TO: Henry James Bicycles Inc / MCI ID: 509-6370
Subject: MCI Mail Problems

Dear MCI Mail CSGnetters:

There have been some problems with the CSGnet-MCI Mail interface.
Hopefullly, these problems have now been solved.

You may have received some messages from the LISTSERV telling you of being
deleted and added to CSGnet. This should end now, at least for the time
being. You are all on CSGnet now using the following addresses:

4912499@MCIMAIL.COM	CLARK Bob, Cincinnati, OH
4742580@MCIMAIL.COM	FORSSELL Dag, Valencia, CA
4972767@MCIMAIL.COM	WILLIAMS Greg, Gravel Switch, KY
5096370@MCIMAIL.COM	FOLSON Hank, Henry James Bicycles

You may have missed CSGnet messages (all or some) from last weekend to
today. All these messages are contained in the CSG-L LOG9208B archive. To
get a copy of this file, send the following message to
LISTSERV@VMD.CSO.UIUC.EDU

GET CSG-L LOG9208B

You will probably also receives full copies of all messages you send to
CSGnet, at least until I can get MCI Mail to remove the leading zeros that
it puts on your addresses in the return address field.

I am sorry for any inconvenience this may have caused.--Gary

Date: Wed Aug 12, 1992 10:16 am PST

Subject: for Wayne

Wayne,
Your "The Synergy of Involuntary and Voluntary Action" is perfect for my Causation paper! i can't believe I didn't look at it before--i wish I had. Thanks!

Mark Olson

Educational Psychology 210
College of Education
Univ of Illinois at Urbana-Champaign
phone: (home) 351-8257
(office) 244-8080

USmail: 405 South 6th St. #4
Champaign, IL 61820
e-mail: (Internet) m-olson@uiuc.edu
(Bitnet) FREE0850@uiucvmd

Date: Wed Aug 12, 1992 12:36 pm PST
Subject: Imitation vs Modelling

[From Rick Marken (920812)]

Oded Maler (920811) says:

> But for engineers (unlike reverse-engineers) this
>problem does not exist, because they build the boxes, they know the
>I-O ("S-R") characteristics of (some of) the components they put in
>the loop, and all your protests against S-R, which might be valid in
>the context of psychological experiments, seem like fighting
>non-existent windmills. They don't know and don't care about ID-IV
>tests and the only arguments you may have with them is whether it is
>useful or not to use the language of hierarchical servo loops,
>reference signals etc. in order to achieve interesting performance (I
>refrain from using the word "behavior"..).

The argument I have with roboticists is the same as my argument with psychologists (both groups, incidentally, believe that I am "fighting non-existent windmills"). Both groups are trying to find a way to imitate behavioral appearances rather than model the underlying processes that cause those appearances. It all comes down to understanding the difference between behaving (what we see people -- or artifacts -- "doing") and controlling (what we can't see people -- or artifacts -- doing, viz. keeping perceptual variables at reference levels using whatever means is demanded by the situation). This is much more than a matter of getting people to "use the language" of PCT. It is a matter of understanding what it means to control. I agree that roboticists often use control systems in their efforts to build machines that put on "interesting performances". What I object to is the orientation towards the performance -- rather than the cause of the performance; the control of perception (at least in living systems). The difference to me is between imitation (the current goal of psychological theorists and roboticists) and modelling (finding the underlying organization that produces the appearance). When you successfully do the latter, the behavior of your system will also imitate that of the real system -- but it will do it for the right underlying reasons (and, thus, will imitate the behavior in new situations in which the "imitation only" device will fail).

But I don't really care whether the roboticists ever understand the difference between imitation and modeling; their goals (as I understand them) are simply not interesting to me. I don't need an artifact running around my house bringing me the paper or mopping up the floor. Robotics is interesting to me only as an implementation of a model of real living systems;robots are of value to me only insofar as they further my understanding of human nature. Whether a roboticist succeeds in making his or her toy behave in some interesting way is of no particular interest to me; whether they do it by just imitating behavior or if they do it by developing a model that accidentally includes some of the principles of control, I don't care. I'm not interested in seeing cute behavior (unless I'm at Disneyland). I care about the underlying principles -- not the superficial appearances -- and I see no evidence that roboticists understand the principles of organization of living systems any better than do psychologists.

While I could care less what roboticists do (except insofar as their creations can serve as existence theorms for a paricular model of behavioral organization), I do care what psychologists do and think because they are presenting a model to the world of "how people work" and people assume that these folks know what they are talking about. The fact that psychology is currently based on the wrong model is not just unfortunate in terms of the progress of our understanding of human nature; it is also unfortunate because wrong models of people lead to wrong ways of dealing with people in the the real world. And this produces REAL life problems for all of us. People typically deal with other people in cause-effect terms (as sr devices or output generators). This leads to a focus on overt behavior (the kind that roboticists want to imitate) and efforts to get that behavior to look the way it "should" look. So rather than focusing on what people are trying to control -- and helping them control it so that the means they select are not so disturbing to others -- people tend to try to solve problems manipulatively -- with the attendant social disasters. These social disasters are just not necessary at all. I want a better world for my children and it won't happen unless people start to understand that their basic assumptions about how people work are FALSE. Force won't work; inflexible system concepts won't work, repression won't work; CONTROL WON'T WORK (in the long run) precisely because people ARE controllers. That's the message I care about; and it's why I think an understanding of PCT is so important. If the roboticists already do understand the the principles of control (as you imply) and are just using different language, and if they understand that these principles apply to people as well as robots, then they would be most helpful allies in our efforts to help psychologists (and lay people) understand the nature of human nature.

Best regards Rick

Date: Wed Aug 12, 1992 2:59 pm PST
Subject: sonja, etc.

[Avery Andrews 920813]

We should probably wait until Bill has spent some time with the Chapman book before saying much about it, but I'll say now that he's perfectly aware of the nature of the hyphenated routine names in the description of Sonja.

One of the motives of the work was to produce a complete system where notions of meaning and references were grounded in actual activity, rather than being just empty `semantic markerese' (typeface-change semantics). I think anybody who studies the book will quickly perceive that Chapman doesn't miss much.

As to why he doesn't emphasise the role of control, I think it's quite possible that he doesn't know about the messes that psychology has gotten itself into, & wouldn't care if he did.

Avery Andrews

Date: Wed Aug 12, 1992 11:38 pm PST
Subject: On Blocks world and "real" world

[From Oded Maler 920810]

(Concerning the recent exchange on simplifying assumtion, disturbance-free worlds, modeling etc.)

[In this short unfocused essay I use often the term "Blocks world" which is a common metaphor characterizing symbolic AI toy worlds, where the world consists of objects,(block-a, block-b), properties (green(block-a)), relations (on(block-a,block-b)) and operators that transforms the state of the systems (e.g., remove one block from another). This gives the flavor of the world on which classical AI problem-solving took place.]

****

Suppose for the moment that the object of study of AI, AL, etc. is not defined with any reference to any reality of living systems. It is the study of mathematical platonic entities obeying some rules, interacting with some environment which is also an abstract object with its own formal dynamics. I'm sure that at this level no one will object the validity of various results obtained there (only their relevance to something interesting). I'll give two concrete examples one in AI flavor and one in AL:

1) An algorithm, that given a "goal" (in this setting, a representation of some configuration of the "environment") produces a sequence of "actions" that leads from any state to this state.

2) A set of simple local dynamical rules, that when employed by many similar abstract entities leads to an interesting emergent global "behavior" (=trajectories in the state-space).

Now, we try to see whether there is any relevance to the "real" thing. Considering the first example, we all know (by now) that the world does not come with labels "a" and "b", and making sense of the sensory input is a big mess. We also know that any abstract action that we want to perform should be realized in a smaller scale with disturbances etc. In other words, the real world is not the Blocks-world. On the other hand, we feel from introspection that *some* of our mental activity is done in this (graph-searching, if-then-else) mode, and *if* the lower-level systems take responsibility for all the messy details, uncertainties, and disturbance, it can transduce the real-world into a Blocks-world, and then a higher-level system can operate at this abstract level, and thus the achievements of systems of type (1) above can become relevant. This was (and probably is) the tacit assumption of symbolic AI, cognitive psychology, etc. and without it all this research is rather irrelevant to human mentality and of little relevance to engineering (except for some non-robotics AI such as [non-video] game playing).

There might be two objections to this view:

I) It is impossible to construct such a real-to-abstract transducer without solving the unsolvable categorization problem (and certainly not in real-time).

II) - And this is the most intriguing fact for me - *EVEN* if such a translation exists, except for some small toy problems, the planning problem in such a world is computationally intractable.

All this means that the blocks world is uncomparable in its complexity to the real world (as perceived by our problem-solving level): in one sense it is simpler because its availiability to us subsumes a non-existent unrealistic transducer, and on the other hand it is more complex because many problems that we can solve with our slow brains, can not (provably) be solved in the blocks world using the fastest computers.

The second observation might suggest that concerning negative results, simplifying assumption are useful bacause it demonstrates that *even* in a disturbance-free environment, something cannot work - so it certainly won't work when ideal transducers are replaced by noisy ones.

I've been attracted to the bottom-up approach (e.g., Powers, Brooks) hoping that some principles discovered while going up in the levels will explain how, after all, it is possible to plan and solve problems in the real-world in the higher level. In what way does the fact that the symbols on which the higher-level systems work, are "grounded" (to use a popular term) in lower-level interaction with our particular concrete world, make problem-solving easier than is some arbitrary abstract world (and this is also, I think, related to some of the ecological ideas of Gibson).

Ok, I see I'll not touch the AL example this time. Feed-back (of any polarity) is welcome.

--Oded

Date: Thu Aug 13, 1992 7:45 am PST
Subject: Keijzer Intro

[from Gary Cziko 920813.1028]

Here is an introduction from a new CSGnetter whom I just put on the network.--Gary

P.S. Bill Powers: Why don't you say hello to Fred and share a few words on how PCT is similar and different from Manturana and autopoeisis (I can never spell that right).--Gary
==

>To: Gary Cziko
>
>From: Fred Keijzer,
> Institute of Experimental and Theoretical Psychology,
> Postbus 9555, 2300 RB Leiden, Netherlands.
> e-mail: KEIJZER@RULFSW.LEIDENUNIV.NL
>
>Recently I discovered the work of W.T. Powers, and I would be very interested
>to learn more about this 'Perceptual Control Theory'. Greg Williams mentioned
>CSGnet as the proper place to be for this purpose, and so I request
>subscription to this mailinglist.
>
>As an introduction:
>
>I'm a theoretical psychologist, working in cognitive science on the problem
of
>fitting the concept of representation into a natural science context. The
>'reference signal' of control theory seems to offer possibilities in this
>direction. I am also interested in any connections between control theory and
>the work of Humberto Maturana, whose ideas are in many ways similar to those
of
>Powers.
>
>Groeten,
>
>Fred Keijzer.

Date: Thu Aug 13, 1992 8:41 am PST
From: Dag Forssell / MCI ID: 474-2580
TO: Gary (Ems)
Subject: MCI mail connection

[From Dag Forssell (920813) direct

Gary:

Thanks for messages and handling of MCI interruption. I downloaded Log 9208B and found that I was cut off already on Saturday morning. Got a rich harvest.

The last message Friday was Avery Andrews (920808). Right?

The listserver refused my command "set CSG-L ack" with the reason that I was not subscriber. This *after* you had put me on again.

If you are controlling for minimum involvement on our part, perhaps you should let me subscribe myself, zeros and all, so that the listserver will recognize me when I address it through MCI mail (in MCI's standard of course, which I cannot influence). Then you can delete me again as your good looking address.

Thanks! Dag

Date: Thu Aug 13, 1992 9:47 am PST
Subject: Your removal from the CSG-L list

Dear subscriber,

As of Thursday, August the 13th of 1992, you have been removed from the CSG-L
distribution list (Control Systems Group Network (CSGnet)) by "CZIKO Gary, U of
Ill., Urbana" <CZIKO@UIUCVMD>.

Virtually, The LISTSERV management

Date: Thu Aug 13, 1992 10:13 am PST
From: g cziko
EMS: INTERNET / MCI ID: 376-5414
MBX: g-cziko@uiuc.edu

TO: * Dag Forssell / MCI ID: 474-2580
TO: Robert K. Clark / MCI ID: 491-2499
TO: Hortideas Publishing / MCI ID: 497-2767
TO: Henry James Bicycles Inc / MCI ID: 509-6370

Dear MCI Mail CSGnetters:

0004742580@MCIMAIL.COM FORSSELL Dag, Valencia, CA
0004912499@MCIMAIL.COM CLARK Bob, Cincinnati, OH
0004972767@MCIMAIL.COM WILLIAMS Greg, Gravel Switch, KY
0005096370@MCIMAIL.COM FOLSON Hank, Henry James Bicycles

After some more fiddling with your addresses, I'm pretty sure I got it right this time (I put one too many leading zeros in last time!). Your address in the CSGnet list is now the same as the address given as the return address in your messages sent to CSGnet and anywhere else via Internet.

This means two things. First, it may be that in the past you received copies back from the network of all messages you sent. This should no longer happen (it's not supposed to). Second, you can now control your status on the network since you will now be recognized as a subscriber. If you don't already receive acknowledgments of messages sent to CSGnet, you should send the following message to LISTSERV@VMD.CSO.UIUC.EDU (DON"T send this to csg-l):

SET CSG-L ACK

(Dag, this should work now for you; let me know if it doesn't)

Other neat things you can do now is temporarily interrupt your reception of messages from CSGnet (e.g., you go on vacation)

SET CSG-L NOMAIL

When you come back to tell LISTSERV

SET CSG-L MAIL

Neat, eh?

Again, I'm sorry for any inconvenience I may have caused you (and even sorrier for all the inconvenience I caused myself).

Your obedient net servant, Gary

Date: Thu Aug 13, 1992 12:55 pm PST
Subject: Blocks world

[From Bill Powers (920813.1400)]

Oded Maler (920810) --

A couple of points concerning your "unfocussed essay."

In defining the formal or abstract approach to AI and AL you give two examples:

>1) An algorithm, that given a "goal" (in this setting, a representation
>of some configuration of the "environment") produces a sequence of
>"actions" that leads from any state to this state.
>
>2) A set of simple local dynamical rules, that when employed by many
>similar abstract entities leads to an interesting emergent global
>"behavior" (=trajectories in the state-space).

Although you say that there is "some environment which is also an abstract object with its own formal dynamics," I don't see in either example any acknowledgement that this environment may include unpredictable disturbances. I know that this isn't supposed to be the real environment, but even in a formal environment, all the essential properties of a real one have to be included if the model is ever to be applied to a real situation.

In (1), the implied model is one that computes the output required to achieve a result in the environment. If the environment contains any unpredictable disturbances that also contribute to the result, there is no "sequence of actions" that will produce a preselected result.

In (2), the same problem exists: there can be no regular state-space trajectories that result from application of a fixed (action) rule unless there are no unpredictable disturbances capable of altering those trajectories directly.

When I speak of disturbances, many people automatically think in terms of adding noise to the input or the output. But the disturbances of which I speak can be, for example, large sustained forces that suddenly appear and persist, that act directly on the outcome independently of action or perception -- for example, a 2-kilogram load suddenly attached to an arm and left there.

Also, there is a tendency to assume that all disturbances are low-level:

>In other words, the real world is not the Blocks-world. On the other
>hand, we feel from introspection that *some* of our mental activity is
>done in this (graph-searching, if-then-else) mode, and *if* the lower-
>level systems take responsibility for all the messy details,
>uncertainties, and disturbance, it can transduce the real-world into a
>Blocks-world, and then a higher-level system can operate at this >abstract level...

Even if the mechanics of behavior are taken care of by lower-level control systems, higher-level perceptions are NOT immune to disturbances. If a cognitive program concludes that the best strategy is to buy 100 shares of IBM at 132-7/8, an open-loop model will simply assume they have been bought and go on from there. In real life, such an order may easily result in buying 0 shares of IBM -- perhaps nobody will sell at that price, perhaps the telephone is out of order, or perhaps your broker informs you that your credit is no good. These aren't mechanical disturbances, but they cause errors in a high-level control system just the same.

It would be possible to program AI or AL programs that let the real world into the loop. These programs would have to check constantly to see whether any commanded act has in fact been carried out. They would have to sense the ACTUAL outcome, compare it with the intended outcome, and treat the difference as the basis for a change in the commands to act. If AI and AL programs were organized like that I would have no more complaints.

Even if it were not true that "the planning problem in such a world is computationally intractable" the presence of unpredictable disturbances would render any plan obsolete before it could be carried out. I don't see the primary difficulty with AI and AL models as being the complexity of the real world. It's the failure to find an organization that can deal with disturbances without having to know about them in advance.

>I've been attracted to the bottom-up approach (e.g., Powers, Brooks)
>hoping that some principles discovered while going up in the levels >will explain how, after all, it is possible to plan and solve problems >in the real-world in the higher level.

I think there's an unwarranted assumption here, which is that planning and solving problems in the abstract actually works. I think that plans and abstract problem-solving very seldom work -- they can't deal with changes in the world that are unexpected or unobserved, and they always leave out critical details. The best that a plan can do is sketch in roughly what might be done assuming nothing surprising happens, and assuming that all necessary information is available. Most plans require emergency modification as soon as they are put into effect. Abstract problem-solving (as in the Prisoner's Dilemma game) usually yields a result only when the game remains abstract and everyone agrees to play by exactly the assumed rules. Try to play it in real life and somebody will cheat (I would confess to the crime in a way that involves a falsehood certain to be discovered).

There are exceptions to the general failure of plans and problem-solving to work. Those are the cases where the plan or solution entails setting up a closed-loop system that takes no outcome for granted, but relies on constant monitoring of results. This is not the approach generally taken.

Rick Marken (9208120 --

Just saying "CONTROL WON'T WORK (in the long run) precisely because people ARE controllers" is enough to convince anyone who already understands the principles of control. To others it just sounds like a slogan. When we talk with people who don't already understand the principles, we have to drop down a level and talk about the REASONS for which control means that controlling people won't work. And so on. That's why I keep harping on disturbances -- they're the main reason that control is necessary, and the main reason why other non-feedback models can't work in the real world. I don't think it's much use to keep saying that control is the answer, when the question to which it's the answer isn't perfectly clear.

Tomorrow I go and get my new 486-33 with supervga screen. What the hell, I won't be eating much when I'm 85 anyway.

Best to all, Bill P.

Date: Thu Aug 13, 1992 2:05 pm PST
Subject: Disturbances/Appearances

[From Rick Marken (920813)]

Bill Powers (920813.1400) says;

>Just saying "CONTROL WON'T WORK (in the long run) precisely because people
>ARE controllers" is enough to convince anyone who already understands the
>principles of control. To others it just sounds like a slogan.

> That's why I keep harping on
>disturbances -- they're the main reason that control is necessary, and the
>main reason why other non-feedback models can't work in the real world.

OK. I was waxing slogany.

The point of my post, however, was to emphasize the distinction between models of appearances (which I called imitations) and models of processes (which I called models). The fact of the matter is that the main reason for this distinction is the existence of disturbances -- in particular, disturbances that are not detectable by the behaving system OR by the modeller of the behaving system. If one does not know that these disturbances are (or could be) affecting the observed results, then s/he will tend to build models that simply produce the results -- ie. that imitate what the modeller sees. But we know (thanks to you) that disturbances (in many different forms) are always having effects on behavioral results. Thus we know that consistently produced results must be controlled -- and what is controlled is a perceptual representation of those results. So we need a control model to produce the results we see. An imitation model won't work because any change in the disturbing conditions would lead to inconsistent results. So imitation models only work in environments where there are no disturbances. Only a control model will produce consistent results when there are disturbances. So, indeed, understanding that disturbances can influence every kind of result that organisms are known to produce consistantly is crucial to acceptance of the control system as the basic organizing principle of the nervous system. And it is crucial to seeing what imitation models are and why they don't really work.

Regards Rick

Date: Thu Aug 13, 1992 6:52 pm PST
Subject: Disturbances

[FROM Dennis Delprato (920813)]

I have been enjoying the discussion of disturbances and simplifying assumptions. For those who wish to talk to non PCT experimental psychologists, I suggest you have an excellent "hook" here.

Start at the beginning. What is the raison d'etre of the much acclaimed experimental method? Control over so-called extraneous variables that are the source of unaccounted-for (error) variance in events of interest. How successful is the IV-DV strategy? Citation, citation, ... Not very, even when statistical control in the form of factorial designs and multivariate statistics is incorporated into the manipulative strategy. Furthermore, aren't we aware of a fundamental problem with highly controlled experimental research, even when (or if) its application permits impressively confirmed predictions? Yes, we find the actual world controlled away such that results are applicable to only a very restricted set of conditions. We end up not knowing any more than before the rigorous experimentation was conduct--with the exception of what applies to an extremely narrow range of conditions. Tied in here are ceteris paribus assumptions in an attempt to salvage the experimenter-as-manipulator methodology.

Thirdly, what is one implicitly assuming when control over extraneous variables is placed in the hands of the experimenter? Aren't we assuming that the individual is pushed and pulled around by external forces, that they are not active participants in their world?

Well, here's a suggestion from an area that is based on the fundamental proposition that the psychologically-active individual always participates fully in all actions. First, consider this notion of extraneous variables; this idea that there are variables that often operate insidiously to befuddle well-intentioned researchers. An alternative is not to deny that something like extraneous variables is found in psychological events, but to re-conceptualize them as disturbances. With the idea of disturbances and the conceptual framework in which they are found as formal constructs, we (a) need not rely on

a wind-weather vane conceptualization of psychological behavior,
(b) can methodologically handle all variables of whatever number
and however classified by the classical experimental model....

Get the point? KISS (Keep It Simple ...)

Basically, I am suggesting a critical look at the classic experimental framework (CEF), a set of assumptions not questioned by adherents to otherwise divergent theories and approaches. CEF-thinking is taught to all psychologists. It is possibly one of the very few points of general agreement in the fractured field of psychology. Not a pleasant thought for those who are not impressed with it. However, it does seem that there may be enough of an undercurrent of dissatisfaction to lead some, at least, to consider an alternative.

Dennis Delprato
Dept. of Psychology
Eastern Michigan University
Ypsilanti, MI 48197 U.S.A.
Psy_Delprato@emunix.emich.edu (Internet)

Date: Fri Aug 14, 1992 2:28 am PST
Subject: High-level disturbances; Ideology

[From Oded Maler 920814]

(Bill Powers (920813.1400)) -

I agree with most of what you say. There's a sub-sub-branch of AI called "Reactive planning" dealing with planning in a non-static world. It seems that *in principle* they might pass you legitimacy tests even though their building blocks are abstract actions. Of course, no one knows yet how to implement it for large-scale problems since servoing in this abstract space is not understood ("If you can't beat it - smash it"..) as we agreed a long ago while discussing "qualitative" control.

I still reiterate the rhetoric question whether for some reason it is harder in a non-grounded abstract space than in an abstract space where the percepts and actions are realized by lower-level ones. And if so, how?

(Rick Marken (920812)) -

Ok. You make your position very clear. So the next time something like Beer's work or another is discussed, your claim should not be that they are wrong but that they are working on non-interesting problems (for you).

Perhaps my knowledge of PCT is so preliminary so I would be less determined and not be so sure which behavior is generated for the "right underlying reasons" and which is not. It is also not self-evident at all that a world where all people know what control is will be a better one - but anyway I'm not interested in Ideology (unless I'm in a historical museum).

Best regards --Oded

Date: Fri Aug 14, 1992 8:33 am PST
Subject: Delprato's paper

[From Bill Powers (920814.1000)] Dennis Delprato (920813) --

When you break your silence you do it in great style. That is an outline for an important paper. I hope you're going to write it.

Here's some more scaffolding for the article (partly continuing the conversation with Oded Maler (920814)). This is part of an outline for something I'm writing, but feel free to use any part of it that seems useful. The following will benefit greatly from some running examples, but I leave them out here for brevity. It contains some experimental usages that may make things clearer -- "disturbances" in particular.

First, break behavior down into actions and outcomes. Most behaviors are named by the outcomes: peeling an apple, toasting some bread, driving to work, splitting some kindling, and so on. The actions that bring about these outcomes aren't normally mentioned.

It's possible to use outcome-names for behaviors because outcomes repeat, even in a variable environment. It's necessary to use outcome-names because outcomes are far more regular than the actions that bring them about. If we described all behaviors in terms of their action components, there would be almost no regularity to name.

Outcomes are affected by an organism's actions and by independent variables that also affect the outcomes. The only way the outcomes can remain regular is for the actions to vary in opposition to the effects of the independent variables, or disturbances.

There are two main kinds of explanation for how actions might be adjusted to achieve the observed regularity of outcomes. One explanation involves compensation. The other involves control. Both require that the organism sense its environment.

The compensation answer requires that the organism sense the states of independent variables that can influence the outcome. This answer has several drawbacks. First, all the causes of disturbances of the outcome must be known; if they are not, compensation isn't possible. Second, the effect of all independent variables on the outcome must be known; that is, when these variables change it's necessary to calculate the effect of the change on the outcome. Third, the action used to counteract the disturbance comes from a different source than the disturbance. So the effect of the action on the outcome must also be known and separately calculated so as to have the right opposing influence on the outcome; in general, the path by which the action affects the outcome is different from the path by which an independent variable disturbs the outcome. Fourth, the links to the outcome from both action and independent variable must have constant properties. If those properties change, to the same extent the calculations will be in error and the outcome will fail to be stabilized. Fifth, the calibration of the sensors and actuators must remain constant.

The control answer requires only that the organism sense the outcome itself. It does not have to sense the independent variables causing disturbances of the outcome. This answer has the drawback, at least for discrete actions and outcomes, that some disturbance must be allowed in order to serve as the basis for the action that prevents further disturbance. However, control will work when the independent variables can't be directly sensed, or when they can't be sensed accurately, or when they vary unpredictably on the same time-scale on which actions take place. Control will also continue to work accurately over large changes in output calibration. Thus control is the more general answer to the question of how consequences are made repeatable in a variable environment.

In some cases both compensation and control can be used. Compensation -- reacting directly to an independent variable capable of disturbing an outcome -- can produce an immediate action that is approximately right for opposing the disturbance. If the remaining disturbance is small enough, a control process based on sensing the outcome itself can succeed in removing the rest of it, achieving exact stabilization of the outcome without any protracted period of error. This combination works best for discrete variables. For continuous variables, for example those involved in balancing upright, compensation does not add any appreciable improvement in performance.

Compensation and control work at about the same speed. Both require sensory computing processes. Both require computation of outputs. A control process, however, is always faster and simpler when any complexity of calculation of output is involved. In a control process the required output is derived continuously from the sensory information, often by a simple subtraction, and varies as the sensed outcome varies. Compensation, however, requires continuous calculation of properties of the environment and the body, which even for simple behaviors can entail calculations of impractical complexity (as in computing the inverse dynamics of an arm in order to move it). Furthermore, the negative feedback involved in control allows high amplification of responses to disturbances, increasing the bandwidth of responses greatly over that obtainable with a compensatory system. In general, a control process is faster than an equivalent compensatory process. And of course, a control process is, in the steady state, far more accurate than a compensatory process, not relying on the environment always having exactly the same properties, not relying on unchanging calibrations of sensors or actuators, and not requiring that the independent variables responsible for variations in the outcome be known at all.

Best, Bill P.

Date: Fri Aug 14, 1992 9:05 am PST
Subject: Disturbances

[From Rick Marken (920813.1000)] Dennis Delprato (920813) --

Great to hear from you. I think your ideas about introducing PCT to non-PCT psychologists are EXCELLENT.

>Basically, I am suggesting a critical look at the classic
>experimental framework (CEF), a set of assumptions not questioned by
>adherents to otherwise divergent theories and approaches. CEF-thinking

>is taught to all psychologists. It is possibly one of the very few
>points of general agreement in the fractured field of psychology.

I agree completely. I've always felt that the difference between conventional and PCT psychology could be best approached from the point of view of methodology (since its the same for all psychologists, regardless of theoretical persuasion). I like your suggestion about moving non-PCT psychologists gently to the concept of control by showing them how they might profitably reconceptualize "extraneous" variables as disturbances. I think you might consider taking the lead on this approach. Perhaps there are some operant conditioning studies where people have felt that the results were a bit too noisy due to these extraneous variables? If you could find a study like this then perhaps you could show that the "noisiness" of the behavior is the organism's efforts to reduce the effect of these variables on a controlled variable. In fact, now that I think of it, the concept of a controlled variable let's you distinguish between "extraneous variables" that are disturbances and extraneous variables that are extraneous. Only variables that have an effect on the controlled variable are disturbances -- ie. the kind of extraneous variable that will actually influence behavior.

Go with this idea Dennis.

Oded Maler (920814)

>It is also not
>self-evident at all that a world where all people know what control is
>will be a better one - but anyway I'm not interested in Ideology
>(unless I'm in a historical museum).

I agree that it is not self-evident. I can see that my sloganizing is coming back to haunt me. What I meant was simply that, if people understood that people were control systems, then they would KNOW for a FACT that their efforts to control people (in order to "make things better") will not work. The reason is simple -- one person's efforts to control another are just disturbances, and will be automatically countered by the controllee. This is not an ideological point -- but a fact about how control systems work.

I have a demo which illustrates this point. It's sort of the inverse of my mindreading program. I know it's a good demo because Gary Cziko liked it. I call it "find mind". Here's how it works:

There are 5 numbers roaming around the computer screen. The behavior of 4 of the numbers is determined by a VERY simple output generation model; at each time interval the x,y specification of position is updated and the number "moves" to the specified screen position. One of the numbers, however, is a control system. The x,y specification at each time interval is the value of the REFERENCE SIGNAL. So this system is controlling it's perceived x,y position relative to a varying reference. The position of ALL FIVE numbers is also influenced by the position of the mouse. If you leave the mouse alone, this just means that a constant is being added to the x,y position of all 5 numbers. When you look at the screen what you see are five numbers drifting around, each taking a different path (determined by the particular waveform that determines its x,y coordinates (or reference signal)). There is NO WAY to tell, by looking at their behavior, which of the 5 numbers is a control system and which are "output generators". The appearance of the behavior of all 5 numbers is EXACTLY the same.

But you can very quickly tell which of the 5 numbers is a control system if you introduce a disturbance by moving the mouse. If you move the mouse abruptly, you will see all 5 numbers shift in the direction of mouse movement, but one number will "bounce back" to its intended path. This bounce is easy to see, even though you don't really know where the number intends to move. So the control system "responds" to the mouse disturbance in a very obvious way; the rest of the numbers show no such bounce -- they just shift their movement to the "offset" caused by the mouse disturbance. The behavior of the control system number is like a "reflex" -- such as the behavior of the pupil when you suddenly shine a light in the eye.

If you apply the disturbance more smoothly then you cannot see it's differential effect on the "control system" number. If you just move the mouse around slowly and aimlessly you will not be able to tell, again, which of the numbers is "alive". But you will quickly learn which number is alive when you try to hold a number at some target position on the screen -- ie. if you try to control the position of a number. If you try to control one of the "output generation" numbers then there is no problem -- you can easily move the mouse appropriately to compensate for the changing x,y position of the number. But if you happen to pick the control system number, then you find that your mouse movements become much more exaggerated. It turns out that the control system is generating outputs to prevent your disturbance (mouse movements) from moving it from its intended path. Your efforts to get that number to stay on target are, in fact, placing you in a conflict with that number. Since I have placed no limit on the amount of output that the control system number can generate, the control system will win the conflict (you will run out of disturbance -- mouse -- first). Some- times, this is the way things work out in real life (like when you try to control little kids) -- but, as I said, not for long, since control systems don't give up.

This little demo illustrates my basic point -- if people are control systems then efforts to control them just create conflict; and conflict is not good for either party because it involves loss of control.

Best regards   Rick

Date: Fri Aug 14, 1992 2:00 pm PST
From: g cziko
          EMS: INTERNET / MCI ID: 376-5414
          MBX: g-cziko@uiuc.edu

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Is it working?

Dag:

Are you now receiving mail from CSGnet? I've been getting some weird returned mail that looks like this:

>Your mail was not delivered to some or all of its
>intended recipients for the following reason(s):

>
>'00047425...' is too long to be a VM userid.

This is the beginning of your address and so I wonder if your address is causing a problem.--Gary

Date: Fri Aug 14, 1992 2:08 pm PST
Subject: Independent Investigator #2?

From Tom Bourbon [14 August 92] Re: Independent Investigator #2?

For the past several months I have been pretty inactive on the net. I was winding down from 25 years at one institution, prior to moving here. We went through the traumas of selling a home of 21 years -- tons of "stuff," a daughter who had known no other home, and the like. Six days ago, we completed the move.

Today, the foundation that provided the largest portion of support for the laboratory here (unexpectedly) cut us off. We have enough money to continue operating the lab at the present level for a couple of months and several proposals are under review, but I am free to look elsewhere. I am looking for honest work anywhere I can find it -- have control handles, will travel. Perhaps PCT will soon have its second "Independent Investigator."

Best wishes,
Tom Bourbon

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

Date: Fri Aug 14, 1992 2:14 pm PST
From: Dag Forssell / MCI ID: 474-2580

TO: list (Ems)
EMS: INTERNET / MCI ID: 376-5414
MBX: LISTSERV@VMD.CSO.UIUC.EDU
Message-Id: 40920814221404/0004742580NA2EM

set CSG-l ack
get csg-l log9208b

Date: Fri Aug 14, 1992 5:18 pm PST
From: Dag Forssell / MCI ID: 474-2580

TO: Garg (Ems)
EMS: INTERNET / MCI ID: 376-5414
MBX: G-CZIKO@UIUC.EDU
Subject: Is it working?
Message-Id: 80920815011808/0004742580NA1EM

From Dag (920814)

Yes indeed. I sent set CSG-L ack, and it came back successful. The Log shows that I did not miss any messages either. So all is well. No error signal here.

Keep up the good work!

I have not yet shopped for another servo, but that will come. Glad we could learn from each other.

Dag

Date: Fri Aug 14, 1992 5:40 pm PST
Subject: what is and isn't

(penni sibun 920814)

> [From Rick Marken (920808)]
>
> penni goes on:
>
> >is this [S-R modeling (rm)] what the
> >alife people claim they are building? it is emphatically not what
> >agre, chapman, brooks and co claim they are building. for these folks
> >and others in ai, for gibsonian psychologists, for
> >ethnomethodologists, etc., agents and their worlds are *not
> >separable*. how can you have a respond to b's stimulus when a and b
> >are the same thing? that's actually not a terribly useful
> >simplification of the argument; try this one: how can you individuate
> >a's-stimulus and b's-response when you can't individuate a and b?
>
> My comments about people building control systems and calling them
> SR devices were based on an article I read in a computer graphics
> magazine where little "bugs" moved around the 2-D surface of the

well, please. you're talking about a different group of people here.

> Trust me. Whatever psychologists call their models, they are SR
> models pure and simple and they are treated this way.

sorry. i was just at the cognitive science conference with several hundred psychologists, and they didn't strike me as s-r modelers; one of the keynote addresses was by turvey, who is a gibsonian (that is, has an interactionist model of organism/environment), another was by bates who talked about recent research in aphasia. i got my undergrad degree in psych at harvard, a bastion of behavorism in the latter's heyday; i was taught behaviorism only as part of the history of the field. if anything, in my experience, psychologists are too likely to be cognitivists and not at all likely to be behaviorists (and a lot of them are too smart to be emphatically either, at least all of the time). most ai-types are raving cognitivists.

on my view, you talk like a cognitivist. partly cause you oppose pct to behaviorism. more tellingly, cause you explicitly locate control in the head, just as cognitivists do. i think it would be a lot more constructive for me to understand your position if, instead of distinguishing your view from behaviorism, you distinguished it from cognitivism. care to give it a shot?

you ignored my radical challenges (repeated above), but they are very much to the point. in order to convince me, and other people (characterized as above) who loosely consider themselves to be interactionists, or situated activity theorists, or whatever, you need to step back from the pro-/anti-behaviorism debate. interactionists *just don't find that an interesting question*.

if you think it's too weird to consider that organism and environment are fundamentally inseparable, why not consider a slightly different take on it: action and perception are inseparable; there is no way you can draw a line and say that on one side is the organism's action and on the other is the organism's perception. on this view, what could it mean to locate control in the head of the organism?

on the interactionist view, control is neither in the head not in the environment.

to conserve my hands, and because they've had years more practice explaining this than i have, i include excerpts from two authors on the ``interactionist'' view. this first bit is from an unpublished manuscript of phil agre:

> A worldview might be decades old, but its rhetoric will continue to encode its formative disputes. Behaviorists, in their very discourse, will be forever locked in battle with introspectionists, and the very discourse of
> cognitivism is still locked in battle with behaviorism. Cognitivism and behaviorism are really two sides of the same coin, sharing a system of metaphors and simply disagreeing on certain important propositions within that common horizon.
> [...]
> In my view, what actually characterizes the cognitivist worldview is [...] a certain system of metaphors. These metaphors begin by marking out a firm distinction in kind between the mental Inside and the world Outside.
> [...]
> Gibson's theory of perception poses a challenge for cognitivism because it locates perception squarely in the evolving relationship between agent and world.
> [...]
> My point is that the metaphor-system of Inside and Outside finds itself unable to make any stable sense of concepts which reside neither in the
> agent nor in the world but precisely in the relationship and interaction between the two.
> [...]
> We wrote Pengi out of dissatisfaction with STRIPS-like planners and with world

models. (Pengi, our program, plays a video game called Pengo.) While not intended as a general architecture, Pengi illustrates certain ideas. It implements a version of Ullman's notion of "visual routines" (1984). [...] Pengi also illustrates some of the structure of an agent's interactions with a structured environment.
[...]
Like Gibson, we locate certain phenomena in the relationship between an agent and its world.
[...]
So where are we? The critical issue is whether one's categories locate things in agents and worlds separately or in the relationship between them.

what follows is the abstract from beth preston's ``behaviorism and mentalism--is there a third alternative?'', to be published in an upcoming issue of _synthese_.

Behaviorism and mentalism are commonly considered to be mutually exclusive and conjunctively exhaustive options for the psychological explanation of behavior. Behaviorism and mentalism do differ in their characterization of inner causes of behavior. However, I argue that they are not mutually exclusive on the grounds that they share important foundational assumptions, two of which are the notion of an inner-outer split and the notion of control. I go on to argue that mentalism and behaviorism are not conjunctively exhaustive either, on the grounds that dropping these common foundational assumptions results in a distinctively different framework for the explanation of behavior. This third alternative, which is briefly described, is a version of non-individualism.

what i'm still trying to figure out about pct is whether the methodology, stripped of the theory, is theory neutral enough that one could build interactionist frobs with it. this is exactly avery's undertaking with trying to recast sonja in pct terms; i find this quite interesting.

cheers.

--penni

Date: Fri Aug 14, 1992 5:46 pm PST
Subject: Re: Re independent investigator #2

[From Dick Robertson]
Ouch! That really smarts. How can they do this to you after you've given up everything? Well, I hope it turns for the best, though it sounds grim at the

Date: Fri Aug 14, 1992 6:04 pm PST
Subject: RE Plooij Papers

[from Dick Robertson]

Hello, Dag. I got your request for Plooij's papers just after our summer term closed, and I'm going off to Wisconsin for a weeks vacation. But I will send them the end of the week after next when school reopens. Best wishes, Dick.

Date: Fri Aug 14, 1992 8:40 pm PST
Subject: PCT vs cognitivism

[From Rick Marken (920814.2200)]

Well, thanks to penni for giving me something to do this evening that's more fun than watching "Lifestyles of the Rich and Famous" (actually, I've never watched it for fear that one of my old high school buddies would be on).

Ok. That's the header. Now to business.

penni sibun (920814) says:

>sorry. i was just at the cognitive science conference with several
>hundred psychologists, and they didn't strike me as s-r modelers;

They've got new words for it. Let's just say they are cause-effect modellers. Actually, lets not even say they are modellers.

>one
>of the keynote addresses was by turvey, who is a gibsonian (that is,
>has an interactionist model of organism/environment)

I won't say a thing.

Well. OK I will refer you to the two papers in chapter 6 of my "Mind Readings" book. They are dedicated to Turvey and his bunch. I don't think he's ever read them, though, unless he was a reviewer. I don't think he would like them.

I must also mention that Turvey has the honor of having published the longest, most systematic and most incorrect critique of PCT to date. If you want to see ignorance posturing with pompous authority, check out the first chapter (pp 1-39) in G. Stelmach (Ed.) Information processing in motor control and learning. New York: Academic Press (1979) -- the article is by Fowler and Turvey. A must read for PCTers.

>i think it would be a lot more
>constructive for me to understand your position if, instead of
>distinguishing your view from behaviorism, you distinguished it from
>cognitivism. care to give it a shot?

Sure. Though I think we've done it before.

There are many incarnations of cognitivism. But what is common to all of them (and what is commonly wrong) is that they view behavior as the end result of mental processing. PCT says that it is perception that is the result of "mental events" (reference signals). This apparently small distinction is as big as the distinction between behaviorism (which sees external events as the cause of behavior) and PCT. There are several cognitive "models" that illustrate my point. One I particularly like is a "hierarchical control" model of movement sequence generation proposed by Rosenbaum. This model assumes that the results of cognitive plans (the little demons moving around the hierarchy) will be exactly as expected. There is no notion that these results are also influenced by external disturbances and that it is, thus, highly unlikely that repeating the "plan" will repeat the result. What is probably happening is that the subjects are controlling (among other things) perceptions of transitions, sequences and configurations and maintaining these perceptions at intended levels.

There are many other examples of cognitive "models" (I use quotes because they never really try to make these models work in real environments; if they did, they would learn a lot, FAST) but they all have one thing in common -- once all the processing has occurred, behavior just "pops out" as the end result. Sorry, but behavior CANNOT happen that way. So, no matter how clever the processing carried out by these "models", they cannot behave; so they cannot be correct.

There are some cognitive models that have some closed loop character- istics; and probably qualify as control system models of the control of higher level variables. An example is Simon's theorm prover (and current models that are similar, like the prolog language itself). These models try to satisfy certain logical conditions (the reference condition) by trying to satisfy other logical conditions in their database. These are control models that work in a logical (not a real) environment (the database). These are interesting models and they may eventually become useful as models of what Powers calls "program level" control. I have no beef with these models -- I think they are very clever. They are just not yet connected in a realistic way to behavior.

> in order to convince me, and other people
>(characterized as above) who loosely consider themselves to be
>interactionists, or situated activity theorists, or whatever, you need
>to step back from the pro-/anti-behaviorism debate.

I'm game.

>if you think it's too weird to consider that organism and environment
>are fundamentally inseparable, why not consider a slightly different
>take on it: action and perception are inseparable; there is no way
>you can draw a line and say that on one side is the organism's action
>and on the other is the organism's perception. on this view, what
>could it mean to locate control in the head of the organism?

In Bill Powers' "Living control systems" page 251-252, precisely this question is answered, simply and quantitatively. The short answer is "because the organism amplifies, the environment dampens". (my quotes, not Bill's).

>on the interactionist view, control is neither in the head not in the
>environment.

Sounds fine; but it turns out to be wrong -- a common problem when trying to understand the world with words instead of models.

I would like to learn about interactionism. But I didn't really learn much from the quotes you provided. Perhaps you could describe an interactionist model of some very simple phenomenon. Matter of fact, perhaps you might start by explaining the nature of the phenomenon that interactionism explains. Does it explain the phenomenon of control? If so, how?

Best regards Rick

Date: Fri Aug 14, 1992 11:11 pm PST
Subject: Re: PCT vs cognitivism

[Martin Taylor 920816 03:00] Almost back from holiday.
(Rick Marken 920814.2200)

>The short
>answer is "because the organism amplifies, the environment dampens".

Not always true, and certainly not true of that part of the environment that contains contro∞ensitive parts of the environment that don't (e.g. bifurcation points, such as things balanced on table edges). Sufficiently often true that it is a beguiling temptation to take it as truth.

Look at my mirror diagram of a few weeks ago to see another view, according to which control is indeed exerciesed in the head but on an environment that cannot be seen as "outside" except by the controller. (There are better versions of the mirror diagram in my abstract of my Paris talk, which I have sent to Bill, and will send if anyone else urgently wants it (50 pages, so I don't want to send too many).

(Lots of interesting stuff been going on in the time since the meeting. I don't suppose I'll comment on much of it (do I hear sighs of relief?), but this just struck me as worth a note.)

Martin

Date: Sat Aug 15, 1992 6:46 am PST
Subject: Separating action & perception

[From Bill Powers (920815.0800)]

Penni Sibun (920814) --

You say some things, and represent others as saying some tings, that fall strangely on my ear. For example:

>for these folks
>and others in ai, for gibsonian psychologists, for
>ethnomethodologists, etc., agents and their worlds are *not
>separable*. how can you have a respond to b's stimulus when a and b
>are the same thing?

If they can't separate agents and their worlds, they're trying to use the same model for the inanimate environment and for living systems. This is simply a mistake: the inanimate world does not run by the same rules as the animate one. Of course if you jack yourself up to a high enough level of abstraction, everything is the same as everything. But then you have nothing interesting to say, because all the interesting things we say about ourselves and the world are about relationships between independent things. If you can't tell things apart, because of having abstracted too far, then you can't see any relationships.

If you (generic) don't distinguish between action and perception, then there will be no place for a physical world that is influenced by action and that is represented in turn in perception. It's hard for me to see how a person who doesn't make this distinction can understand what it means to say that action controls perception, unless that person means that perception IS action, or that perception follows faithfully from action (which it doesn't -- that's not the meaning of "control").

>if you think it's too weird to consider that organism and environment
>are fundamentally inseparable, why not consider a slightly different
>take on it: action and perception are inseparable; there is no way
>you can draw a line and say that on one side is the organism's action
>and on the other is the organism's perception. on this view, what
>could it mean to locate control in the head of the organism?

But action and perception are completely separable; among our most basic and elementary demonstrations of control processes is one that shows an almost total LACK of correlation between a person's actions and the perception that they control. If you don't understand this fundamental separation, then you've missed a basic point about how control works.

This is such a basic concept that I don't think we can understand each other until we reach some sort of agreement on it.

Think of the car and driver. The driver controls a perception of the car's position in its lane by the action of applying a turning-moment to the steering wheel. Essentially all of the driver's actions are present only to oppose the effects of independent disturbances on the car's path. Those disturbances are unpredictable and invisible. Nevertheless, the driver's actions correlate highly (negatively) with their net vector sum. But the driver's actions correlate hardly at all with the visually-detected position of the car on the road, which is the controlled perception.

How does this jibe with saying that action and perception are inseparable?

Best,   Bill P.

Date: Sat Aug 15, 1992 6:55 am PST
Subject: Good Luck, Tom

from Ed Ford (920815:0846)   Tom Bourbon -

What a blow! Keep us in touch and good luck.   Best, Ed

Date: Sat Aug 15, 1992 7:12 am PST
Subject: Amplification: organism vs environment

[From Bill Powers (920815.0930)]

Martin Taylor (920815) --

Rick Marken says:
>>The short
>>answer is "because the organism amplifies, the environment dampens".

and Martin replies

>Not always true, and certainly not true of that part of the environment
>that contains control systems. Not true in sensitive parts of the
>environment that don't (e.g. bifurcation points, such as things >balanced
on table edges). Sufficiently often true that it is a >beguiling temptation
to take it as truth.

When control systems come in contact, that's a rather special case. Generally they avoid doing so in a way that pits their outputs directly against each other. The reason is that they're BOTH high-gain systems. This doesn't negate Rick's statement. The non-controlling environment is not, in general, a high-gain system.

The rest of the environment contains so few instances of significant amplification that they're pretty unusual. It's hard to think of a case in which the gain of a human controller isn't higher than that of the environment by a wide margin. When a car moves along the crown of a cambered road, the car is unstable about a knife-edge line along the crown; its path bifurcates. But a human driver has no problem keeping the car centered on the crown within a small range. The book teetering on the edge of the table is easily kept from falling by a person who lays a finger lightly on it. Even ten-year-olds have no problem with balancing a broomstick on end. To find instances of significant amplification in the physical world, you have to talk about explosions and things that snap from one state to another faster than a human control system can control them between states. There aren't just a heck of a lot of such phenomena. Most of those that do occur are man-made (like explosions and switches).

As to the mirror diagram, it is an epistemological tool. This diagram could be expanded to show the hypothetical world of physical reality between the two mirroring hierarchies. When we analyze control processes, that is essentially what we do: we apply physical analysis to our experiences and attribute them to physical properties of the world between the systems. That's the only practical thing to do. The mirror diagram itself exists only in the perception of the beholder; it's an hypothesis about the relationship between two control hierarchies, as envisioned by one of them. To continue this process of abstraction is to head for a state of mute contemplation. I think we have to admit that it's all hypothetical, and try to formulate our hypotheses in the way that's most tractable both for theorizing and for experimenting. My formulation puts a physical world running by physical laws between the organisms which are control systems running by closed-loop laws. That seems to work passably well. And in that formulation, Rick's comment is right nearly all of the time, if you understand "environment" to mean the milieu between control systems.

Best, Bill P.

Date: Sat Aug 15, 1992 1:28 pm PST

Subject: Mirror diagram

From Dag Forssell (920815) Direct

Glad you are back. The net has been a bit quiet without you.

I am very interested to see how you have made use of your creation.
Please send me a copy of the 50 pages. Send bill if you want.

Thanks!

Dag Forssell
23903 Via Flamenco
Valencia, Ca 91355-2808
Phone (805) 254-1195 Fax (805) 254-7956
Internet: 0004742580@MCIMAIL.COM

Date: Sat Aug 15, 1992 11:01 pm PST

[from Avery Andrews (920815)]
(penni sibun 920814, on marken 920808)

> on my view, you talk like a cognitivist. partly cause you oppose pct
> to behaviorism. more tellingly, cause you explicitly locate control
> in the head, just as cognitivists do. ...

What I'd say to this is that PCT is not cognitivist, because the *phenomenon* of control is only found when an organism (or other `Agent') is put in an appropriate environment (e.g. if the road is covered with an oil-slick, one ceases to observe control of the direction in which the car is going). The phenomenon of control is just what Rick says it is: when you introduce disturbances that you would expect to change some aspect of the environment, but that aspect doesn't change, because the Agent consistently does things that `mysteriously' counteract your attempted disturbances.

There is a terminological difficulty here in that in normal usage, `control' does *not* imply any capacity to cope with unpredictable disturbances as they arise (I have retained a copy of one of Randy Beer's postings that documents this point extensively, which I could send to anyone who'se interested in looking at it). So for example, most people would call the spikey cylinder in a music box its `control unit', but this object does not effect control in the PCT sense, since it doesn't counteract disturbances in any systematic way.

PCT goes on to say (at least) two further things:

a) an Agent that is efficacious in the real world must effect control

b) the only way control can be effected is by means of certain kinds of internal arrangments - perceptors, comparators & effectors appropriately connected so as to constitute negative feedback loops in the context of the actual environment.

The argument for (a) is that the sorts of things that serious Agents might need or be designed to achieve are going to be subject to constant and unpredictable disturbances, which will then need to be counteracted as they arise, and the argument for (b) is that no alternative arrangments that effect control have been proposed (for example, the reason that Sonja could cope with conveyerbelts, whirlwinds, and, let us say, crunchbirds that would occasionally eat her monster-killing missiles in mid-flight, is that her central architecture is full of little control systems, even if Chapman didn't think of them that way).

So I'd see PCT as a subtype of interactive AI, which makes some additional claims about how to make sense out of how the internal structure of an Agent is related to its `behavior' (e.g. what happens in environments that contain it). I'd take what I've come up with from cogitating about Sonja so far as supportive of these claims (and am encouraged that Penni finds it interesting) but there's still a lot I don't know, & there might be all sorts of genuinely efficacious but not-PCT organization in Sonja that I have missed. But my experience so far is that PCT ideas really are helpful in figuring these things out (recall that I got on to Chapman & Agre when I showed Penni my proposed control system organization for getting beer, & she said that it was like their stuff).

What I would like to do is design and/or analyse some relatively large-scale agent along PCT lines - the crowd people and astro are a start, but they are obviously too rudimentary to convince people of the utility of the approach, in part because they don't have any `tactics'. Crowd people, for example, avoid obstacles, but they don't really navigate in the way that Sonja does. if you set things up so that there is a wall of stationary people between a mobile crowd person and her goal:

```
                  P
                  P
     G            P                            Goal
                  P
                  P
                  P
                  P
```

(you can position stationary people by making them `active', but setting their gains to zero), G will head straight for the Goal until she is about to crash into the row of P's, and will then veer off one way or the other, not consistently taking the shortest way. But Sonja (or my dog) would head immediately to whichever end of the wall would give the shortest path around. I conjecture that the way to get the crowd people to navigate would be to give them more sophisticated perceptual functions (maybe even tuning the ones they have would do the trick - tho I haven't managed to do this), but to say this is not of course the same as to actually do it (Sonja in effect controls for not having a wall between where she is and where she wants to be going, according to my current understanding).

Avery.Andrews@anu.edu.audi

Date: Sun Aug 16, 1992 12:20 am PST
Subject: Re: Amplification: organism vs environment

[Martin Taylor 920816 01:30] (Bill Powers 920815.0930)
>
>The rest of the environment contains so few instances of significant
>amplification that they're pretty unusual. It's hard to think of a case in
>which the gain of a human controller isn't higher than that of the
>environment by a wide margin.

I don't think the point is worth belabouring, but elasticity has much the same perceptual effect as an interaction over a physical variable with another controller that has a different reference. At one extreme we have a totally non-resisting medium, in which (apart from disturbances) the physical variable is readily induced to provide the desired percept. At the other end of the scale, we have crystalline elasticity, in which the controller can have only a trivial effect on its percept, because the world resists strongly any attempt to deform it. In between, there is a whole range of impedances, both resistive and reactive. Whatever the real-world impedance might be, it affects the dynamics of control, and the behaviour of the controller is intimately bound up with the real-world impedance of the provider of the percept. So it is difficult to say straight out that the control is totally in the head, even when there is no identifiable active controller affecting the same physical complex variable. (Impedance is a term from linear dynamics, but consider any nonlinear extension of it, in the above).

>As to the mirror diagram, it is an epistemological tool. This diagram could
>be expanded to show the hypothetical world of physical reality between the
>two mirroring hierarchies.

My reference to the mirror diagram was not to the one linking two controllers in the Layered Protocol structure, but to the one depicting the idiosyncratic structure of the Universe as seen by any one controller. (Bill, did you ever get my Paris abstract with these pictures in it?).

Both the mirror diagram and the considerations of loop dynamics as affected by "real" world impedance argue that there is validity to the view that the internal and external worlds cannot be completely separated. All we have to work with is perception, after all.

>My formulation puts a physical world
>running by physical laws between the organisms which are control systems
>running by closed-loop laws. That seems to work passably well.

Yes, mine, too, in the Layered Protocol structure as integrated with PCT. And before, but not so explicitly.

Martin

Date: Sun Aug 16, 1992 6:17 am PST
Subject: Re: amplification, organism vs environment

[From Bill Powers (920816.0800)]

Martin Taylor (920816.0130) --

[I haven't received your "Paris paper" -- I think I would have noticed a 50-page package!]

You say

>... but elasticity has much the same perceptual effect as an interaction over
>a physical variable with another controller that has a different reference.
At >one extreme we have a totally non-resisting medium, in which (apart from
>disturbances) the physical variable is readily induced to provide the desired
>percept. At the other end of the scale, we have crystalline elasticity, in
>which the controller can have only a trivial effect on its percept, because
the >world resists strongly any attempt to deform it. In between, there is a
whole >range of impedances, both resistive and reactive.

I think we're talking at cross-purposes here. I was speaking about amplification, an input-output property, while you're talking about resistance to disturbance, a possibly closed-loop property. Amplification (more properly, power gain, which takes impedances into account) entails a source of energy, a power supply, that can be tapped so that a variation in the output signal contains a greater energy flux than a variation in the input signal does. This means that the output of an amplifier does more work than is done on the input of the amplifier. If input and output impedances are equal and resistive, the output amplitude is greater than the input amplitude (unless the loop is closed). This is an aspect of "dissipative systems" that doesn't seem to get mentioned much.

When a control system's output meets with "crystalline" resistance, its power output is dissipated mostly internally. The deflection that does result (assuming that's the sensed and controlled variable) is small, so the loop gain of the control system is smaller than it would be for a more compliant and more easily-sensed environment. The internal power gain of the control system, however, is still large: a small deflection creates a small change in perception, involving a few microergs of energy transfer. The resulting output change may involve many pounds of force, acting through the distance of the deflection. Even taking impedances into account, there is a large power gain through the control system. The open- loop change in force per unit deflection is larger than the open-loop change in deflection per unit force.

The energy put into the crystalline substance compresses it, storing at most an amount of energy equal to the force times the deflection. The largest amount of energy that can come out of the crystalline substance (pushing back against the disturbance) is the amount put into it. So the power gain of the non-living substance is at most unity.

Only when some external energy source has put some energy into a non-controlling part of the environment -- lifting the broomstick and standing it on end -- can there be any power gain in the non-controlling environment. A small force applied to the upright broomstick is amplified as the broomstick leans and begins to lose its stored energy. The end of the broomstick can, if its motion is resisted, exert a greater and greater force on something else as the angle departs from vertical, producing much more power than the initial push put into it. So for the moment at least, the broomstick can display power amplification.

However, the power amplification of the control system balancing the broomstick can easily put enough energy back into the broomstick to raise it to the vertical again. The optical energy taken from the broomstick by looking at it is negligible; the muscle energy put back into the broomstick by moving its bottom end is immensely greater. So the control system's power amplification is incomparably greater than that of the broomstick.

Even if the control system were trying to dent the surface of a crystal by pushing on it, the same situation would apply. If we supplied a microscope that magnified the perception of the amount of dent, there would be a range of deflections over which control would be excellent -- we've raised the loop gain by using the microscope, which can be thought of as an impedance-matching device. The power gain has not been changed, but now the control system can sense the deflections more accurately and vary its output force to create any deflection in the possible range. The possible range, of course, is set by the maximum available output force, again a matter of impedance matching. If the output is applied to a compound lever minifying the output movements, which does not change the power transfer but does lower the output impedance, the amount of deflection can now be as large as the control system wants it to be, up to the point where the crystal fractures.

I used this principle in fine-controlling the position of the 300-pound carriage of a grating ruling engine. The input (sensing) impedance was matched by using a laser interferometer to measure the position of the carriage; the output device was a voice coil impedance-matched through a compound lever to bend part of the steel casting. As the required deflections were only about 0.5 wavelength of light, the voice coil thrust, multiplied about 10,000 times by the compound lever, was quite sufficient to achieve the desired deflections. The power gain of this system was probably on the order of a million, after all losses are taken into account. The amplitude loop gain was about 1000. The open-loop power gain would have been -- God knows. Billions and Billions, to quote a famous person.

I used a similar principle to stabilize the laser, an ordinary $100 helium-neon laser. The controlled variable was the audio-frequency beat frequency between the two modes of oscillation of the laser, and the output function was a ring of resistors that heated the Pyrex tube of the laser.

Once you've played with control on such a microscopic scale, there ain't no such thing as a "crystalline substance" no more. The world is made of soft rubber. When you pushed with one finger on that 7-ton engine, just about anywhere, you could see the output meter swing a little as the control system adjusted the carriage position to resist the disturbance. And of course the engine-room was temperature controlled; just walking into the room could make the meter showing heater power swing full scale. The whole system was a clutter of control systems. Of course you couldn't see any of them working: nothing was happening but the ruling strokes. That was the idea. Absolutely nothing was supposed to happen except that the ruling diamond moved to the right place and stayed there during the stroke. We maintained about 0.005 wavelength accuracy for 30 days at a time. The system is still running, although since I left the rat's nest circuitry has mostly been replaced by a computer. My control systems tended to look more like the nervous system than a printed circuit (a mess). My guiding principle for design is called "reorganization."

Best, Bill P.

Date: Sun Aug 16, 1992 8:17 am PST
Subject: Welcome, Fred

[From Bill Powers (920816.1000)] Fred Keijser (920813) --

Hello, Fred, and welcome to our conversation.

I have had some contact with Maturana; in some ways we have similar views, but we diverge greatly in other ways. Maturana's concept of autopoiesis is not worked out in any detail -- it's just a statement that some systems are "self-making." What such a system would have to be able to do in order to make itself isn't explained.

There's a discussion of my views on Maturana in Living Control Systems II, available from Greg Williams.

My biggest argument with Maturana regards his "closed" nervous system. If you're going to assume a model of a nervous system, this entails the whole theoretical structure that is behind it: neurology, physics, chemistry, and so on. So to claim that the nervous system is structurally closed is only a half-baked epistemology. If you admit of a physical nervous system, then to finish baking the idea you have to include the relationships of the nervous system to the physical external world -- which Maturana refers to as the "medium" and treats as an unknown black box. In my opinion he doesn't have a workable model -- only a sort of metaphor.

As to "representation," PCT addresses it in two ways. There is perception, which is a signal-based representation of aspects of the current external world. There are also reference signals, which are signal-based representations of the intended state of the perceptions. The perceptual signals arise from sensory interaction with the external world. The reference signals are internally generated. Which aspect of representation are you concerned with?

Again, welcome.

Bill Powers

Date: Sun Aug 16, 1992 1:54 pm PST
From: g cziko
EMS: INTERNET / MCI ID: 376-5414
MBX: g-cziko@uiuc.edu

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Re: Is it working?

Dag:

>Yes indeed. I sent set CSG-L ack, and it came back successful. The Log shows
>that I did not miss any messages either. So all is well. No error signal
here.

No error signal there means no error signal here. That's good.

>I have not yet shopped for another servo, but that will come. Glad we could
>learn from each other.

I just received my servo cycle. Now to find the time to hook it up!--Gary

Date: Sun Aug 16, 1992 2:03 pm PST
Subject: PCT & Cognitivism

[From Rick Marken (920816)]

I want to thank Bill Powers and Avery Andrews for both,in their own ways, making this a very pleasant weekend for me (despite the fact that the temperature is in the high 90s -- I've enjoyed, however, thinking about the fact that it's winter in Australia; most refreshing).

Avery Andrews (920815) --

That was an EXCELLENT post.

You make several excellent points. First, you correctly point out that the PCT use of the term "control" differs from its use in many other contexts, where it is used as a synonym for "cause" or "determine". I get so used to using "control" in the PCT sense that I forget that it has these other connotations. Your point is very well taken (and put).

Since I'm currently listening to the third movement of a delightful Mozart sonata, I'll even accept your view that you:

> see PCT as a subtype of interactive AI

as long as we can agree that, for living systems at least, it is the correct (or currently best) subtype.

> I conjecture that the way to get the crowd people to navigate
>would be to give them more sophisticated perceptual functions

Yes! Right now the crowd people "adjust" by basically reorganizing (I think there is a bit of random jitter in those higher order references); that's why they rarely get stuck in traps or keep pressing against the same spot on the wall (as in your example). If you want a crowd element to get around obstables more efficiently (which means, control a higher order variable rather than reorganize)
then you have to figure out what they might control in order to properly reset the lower order variables.

But remember, when you build intelligence into the model based on you OWN understanding of the environment, you might be putting in a capability that would not work in other types of environments (and would have to be reorganized out -- that's probably why re- organzation occurs in real systems). What I mean is; if you put in a control system that let's your crowd person efficiently skirt a wall perpendicular to the path of approach, the system may not work for other kinds of wall (or walls with doors). This is just off the top of my head -- but you get the gist?

But you're right; making a "cleverer" control system probably means finding a higher order perceptual variable (like some measure of rate of change of a variable?) that can be used to adjust the reference for another perception so that the little crowd person can get around the wall more efficiently. Good exercise. Of course, the real first step is to see what perceptions efficient wall skirters (like your dog) actually control when skirting walls. That would take some real ingenuity -- testing for your "wall skirter's" controlled variables. Once you know what the dog is controlling, then it's a piece of cake (sort of) to build the model.

Best regards Rick

Date: Sun Aug 16, 1992 2:34 pm PST
Subject: Generating Output

[from Gary Cziko 920816.1715]

Rick Marken (920814.2200) said:

>There are many incarnations of cognitivism. But what is common to all
>of them (and what is commonly wrong) is that they view behavior as the
>end result of mental processing.

and

>There are many other examples of cognitive "models" . . . but they
>all have one thing in common -- once all the processing
>has occurred, behavior just "pops out" as the end result. Sorry,
>but behavior CANNOT happen that way. So, no matter how clever the
>processing carried out by these "models", they cannot behave; so they
>cannot be correct.

This reminds me of a conversation I had with Greg Williams at the last CSG meeting in Durango. Greg was telling me that the arm demo does generate output. It may not calculate trajectories, but it nonetheless generates a behavior based on the error signal of the difference between its reference level and its perception. And the digital nature of computers requires that these calculations be done in discrete steps (in the sequence perception, comparison, action).

When I show Demo2 to someone with computer savvy, he or she can say that what the model does is a reiteration of perception-behavior-perception-behavior, etc. a la TOTE. I realize that if we were all using analog computers this TOTE interpretation could not be made, but by using digital computers don't the demos show in fact that control can be accomplished by generating output based on input (sounds pretty S-R to me)? I would very much like to have Marken and Powers comment on this with critical reflection by Williams. --Gary

Date: Sun Aug 16, 1992 2:58 pm PST
Subject: Why control

[From Bill Powers (920816.1530)] Avery Andrews (920816) --

You make a nice case for PCT, but it could be carried further. This:

>What I'd say to this is that PCT is not cognitivist, because the *phenomenon*
>of control is only found when an organism (or other `Agent')is put in an
>appropriate environment (e.g. if the road is covered with an oil-slick, one
>ceases to observe control of the direction in which the car is going).

... could be expanded by contrast with the environments that Pengi and Sonja are given. The whole digital approach to imitating behavior (Rick's meaning of Imitating) is based on the expectation that there is a reliable causal chain all the way from the top command level to the final environmental result. Peripheral devices used in digital control systems are designed like the computer itself -- an output pulse causes a stepper motor to execute exactly 7.5 degrees (or whatever) of rotation, and the gear or screw on the shaft causes the load to advance or retract by an exactly known and repeatable amount. On the input end, you type a key on a keyboard or receive a modem code, and never get anything partway between A and B. Any error whatsoever is likely to be fatal. Sonja uses a joystick that always produces a motion in one of eight directions -- exactly. When Pengi kicks an ice-cube, the ice-cube always goes along the lines of a grid. And so on. This is an environment in which the lowest levels of action are absolutely precise, predictable, and undisturbed. That's the ONLY reason that these programs can work.

If you substituted a normal continuously-variable environment with continuously-variable disturbances (even moderate and slow ones), and if the means of action were also made continuously variable, there would be no way to reduce the production of action to if-then rules; do it or don't do it. I think this would scuttle most AI and AL programs.

When I talk about disturbances, I don't mean big dramatic things like an oil-spill on the road. I mean just tiny little things: a gust of air, a pebble that rolls a bit underfoot, a muscle that gets a little tired, a slightly awkward stance when you aim, a little overshoot when you whip around a corner, a bit of play in the steering gear. Most action is a succession of events that begin where the previous event left off; we move not in terms of position but in terms of accelerations and velocities which integrate to positions. Every move we make is slightly disturbed, a little different from what was intended, a little off the mark. If it weren't for continuous control operating all the time at many levels of organization, these errors would be cumulative. We would blunder around bumping into doorframes, stumbling over our own feet, knocking glasses of water over instead of picking them up, driving off the road and across the fields, and crushing or dropping eggs when we tried to pick them up. The only thing that allows behavior to work at all is control, even in a quite normal environment.

And control is not a matter of commanding outputs, but of requesting inputs. The outputs vary according to whatever disturbances are happening, regardless of the goal. Only the closed-loop control organization can do this.

>There is a terminological difficulty here in that in normal usage,
>`control' does *not* imply any capacity to cope with unpredictable
>disturbances as they arise (I have retained a copy of one of Randy

>Beer's postings that documents this point extensively, which I could
>send to anyone who'se interested in looking at it).

This terminological difficulty arose because people (including Randy Beer) don't understand how control works. In most usages of this term, the user assumes that setting up the "controlling" conditions is enough to assure the "controlled" result. The role of the controlling entity in the background is overlooked. When we speak of the "tone control" on a radio, we mean that this knob can set the tone to any state between full bass and full treble -- provided that there is a person using it for that purpose, relative to a desired tone. Left to itself, the tone control knob can't control anything. If the speaker is muffled by a towel, the tone control won't turn up the treble. When it's said that temperature controls the rate of a chemical reaction, the same thing applies; the rate of the reaction is controlled by someone's varying the temperature as required to produce a given reaction rate. If the reaction rate varies because of changes in the chemistry, the temperature, left to itself, won't do anything about it.

All such uses of "control" carry the unspoken addition: "all else being equal." But this is a false interpretation of control; control does not require all else to be equal. When control works -- when you make the rate of the reaction constant by adjusting the temperature -- the rate of the reaction is very reliably controlled, even if you vary the temperature in order to keep the reaction rate constant and even if external factors that affect reaction rate require doing this.

Because people don't understand control, they think that simply adjusting the MEANS of control is sufficient to produce a reliable result. They believe this in part because when they see the same result being repeated or maintained, they assume that "all else" MUST have been equal, when in fact it was not. They see control taking place, and don't realize that sensing the result is absolutely essential to success. They don't realize that the adjustments taking place are necessary because of ongoing disturbances. They see someone set the knob to 20 and think that in the future, all they have to do is set the knob to 20 again to get the same result again. Even when they do the controlling themselves, they are unaware of the role their own perceptions are playing.

I've told this before, but it's worth a repeat in this context. A psychologist once insisted to me that to drive a car around a bend in the road, all he did was turn the wheel by the right amount, and the car just went around the curve. He didn't see any reason to watch the road during the turn. I don't know if he's still alive.

When people talk about controlling something by some means, they mean that the something is reliably produced. They don't mean that the controlling circumstances are established and the result varies all over the place. They just aren't aware that in most circumstances, if there were no feedback control involved, setting up the controlling circumstances and leaving them set WOULD allow the result to vary all over the place. Even when setting a control will produce a reliable result for some period of time, they aren't aware of that little feedback loop that works while the control is being set. They aren't watching the controlling device, but the result -- that tells them when the adjustment is "right". They feel the bathwater while turning the hot and cold water knobs. They watch the flame while they adjust the burner on the stove. They feel the resistance build up as they pull the hand brake to "on." But when they describe these control processes, they describe the means: the turning of the faucet, the turning of the knob on the stove, the pulling on the hand brake.

I claim that in just about every case where "control" is used in a way different from that in PCT, the person using the term really expects true control to happen. The outcome is expected to come to a particular state and stay there. But the action or situation that is spoken of as controlling is hardly ever sufficient to produce that result. What is usually meant is "influence" or "affect" or (all else being equal) "determine" . That is not sufficient to produce the expected result of constancy or stability, save in a highly artificial environment constructed to exclude all disturbances and allow only discrete settings of the "controlling" element. And even then, someone must find the adjustment that produces the desired outcome, not by watching the adjustment but by watching the outcome.

Best Bill P.

Date: Sun Aug 16, 1992 4:19 pm PST

A dumb question about Borland graphics in CGA mode -- is there any fast & sensible way to erase circles in CGA Hi Res mode?

Avery.Andrews@anu.edu.au

Date: Mon Aug 17, 1992 7:05 am PST
Subject: Erasing circles

[From Bill Powers (920817.0900)]

Avery Andrews (920816) --

>.. is there any fast & sensible way to erase circles in CGA Hi Res mode?

No. The method is neither fast nor sensible but it's inevitable. It's the method you've been hoping to avoid. You have to erase every pixel by changing it back to the background color.

There are two methods in general:

1. Change color to black (setcolor(BLACK)) and draw the figure in the old position, then change color back to white (setcolor(WHITE)) and draw it in the new position. This usually means saving the old arguments for each instance of a figure, using them with the color set to BLACK to erase all the old figures, changing the saved arguments to the new arguments, then with the color set to WHITE drawing the figure again.

2. Change the mode of writing to XOR. This is done with setwritemode(1). (The normal mode is 0)

In the XOR mode, drawing a figure puts it on the screen, and drawing it again in the same place erases it: it toggles on and off. This is nice because it doesn't erase intersecting lines -- just reverses the common pixels, then on erasure reverses them back.

This is a little trickier because you have to make sure to draw the figure before it's erased (it's erased by drawing again to the same place). Generally you want to erase then write very quickly to avoid flicker, so this means doing a preliminary write of all the figures before entering the loop where you erase before writing. It's slightly faster than the first method, and as mentioned doesn't destroy overlapping figures.

XOR will work for circles, rectangles, lines, polygons (drawpoly), and ellipses. XOR does not work for putpixel, which takes a color argument as its third argument.

NEITHER method works for erasing graphics text! When you write with graphics text colored black, nothing happens to whatever was there. Black is evidently treated as transparent if the background color is already black. I eventually discovered an easy way to do it:

```
setfillstyle(0,0)
bar(x,y,x+textwidth("string to be erased"), y+textheight('W'));
```

This is for top left justified text (settextjustify).

Best, Bill P.

Date: Mon Aug 17, 1992 8:27 am PST
Subject: Re: Separating action & perception

First, I would like to thank you for using a simple example to demonstrate your point. These days it seems that more complicated examples are being used to try to convey a simple or fundamental thought.

Your example is as follows:

> Think of the car and driver. The driver controls a perception of the car's position in its lane by the action of applying a turning-moment to the steering wheel. Essentially all of the driver's actions are present only
to
> oppose the effects of independent disturbances on the car's path. Those disturbances are unpredictable and invisible. Nevertheless, the driver's actions correlate highly (negatively) with their net vector sum. But the driver's actions correlate hardly at all with the visually-detected position of the car on the road, which is the controlled perception.

To me it would seem that drivers do in fact see the disturbances. It is true that they do not see each individual one but they do see the net effect of them by noting the position of the car in the lane (this includes any steering errors also). The action of steering the car so that it stays in the lane seems like a negative feedback type of action solely controlled by sight (excepting for large disturbances that change the wheel position and the change is "felt" and reacted to).

I may be looking at the example employing too much or too little abstractness. I know what you are trying to explain in the example, I am not sure that the example will work but I can not think of a better one myself. My problem is with the driver not being able to "see" the disturbances and not controlling the position of the car on the road. Could you clarify the point a little further it would help me believe in your methods and theories (I have not yet recieved my copy of your book from my University).

salvo, -John
-Student of electron behavior in materials

John van Loon: John_Van_Loon.XRCC%xerox.com@uunet.ca

Date: Mon Aug 17, 1992 10:18 am PST
Subject: Generating Output

[From Rick Marken (920817.1100)]

Gary Cziko (920816.1715) says:

> Greg was telling me that the arm demo does generate
>output. It may not calculate trajectories, but it nonetheless generates a
>behavior based on the error signal of the difference between its reference
>level and its perception.

>When I show Demo2 to someone with computer savvy, he or she can say that
>what the model does is a reiteration of
>perception-behavior-perception-behavior, etc. a la TOTE.

>but by using digital computers don't the demos show in fact that
>control can be accomplished by generating output based on input (sounds
>pretty S-R to me)?

I think it is fair to think of components of the loop in s-r terms. External variables get transformed into perceptions so p = f(i) -- a functional relationship between a stimulus and a response; perceptions get transformed into errors, errors (as you note) get transformed into outputs and outputs get transformed into inputs, completing the circle. As you note, the computer computes these s-r components of the loop sequentially; but it must mimic the dynamic aspect of these functions (the fact, for example, that perception does not instantaeously become p when the input is i) and it does this by only computing a fraction of the computed "response" on each iteration of the sequence and integrating each computed fraction over iterations. Because things work in this nice, dynamic way, it means that the response to a particular input is never exactly the same -- it depends on what went before and what else is currently affecting the variables. So its not really correct to say that the control loop works (ie. controls) by generating precisely computed outputs, although all the components of the loop do produce outputs in response to inputs -- at least they "start" to produce those outputs -- so you can call them sr components. For example, suppose that the s-r relationship between error and output is o = 10*e. So when e = .1, o should be 1; but what really happens is o just starts to change towards 1. So o might be .2 when e = .1. At some later time, as the control system works to get the perception under control, e (which is always changing) might be .1 again. At that point, o might already be .9 so when o starts "responding" to the error, o ends up as 1 -- what it is "supposed to be". The point is that the dynamics of the loop turn the apparatly exact functional mapping of error in output (or of perception into output) into a continuously changing relationship.

But it's probably not worth trying to teach your listener differential equations in order to get the point across that control systems do not work by computing functionally exact outputs (at each point in the loop). The fact of the matter is that when you connect s-r components together in a dynamically stable negative feedback loop, what that loop does is keep its perceptual signal equal to its reference signal (if there is an explicit reference signal) or to zero if there is no explicit reference. This is the fact that was missed by those who proposed the TOTE loop as a model of purposeful behavior. A properly functioning TOTE loop is a control loop; and what it controls is its perceptual variable. The implications of that simple observation are enormous, but only if you notice that that's how TOTE loops work (incidentally, TOTE is probably a poor name since control loops don't EXIT -- unless they are explicitly shut off by another system; so the TOTE loop, when it works (ie. is a control loop) is really a TOT loop).

If there were a simple way to convince your computer savvy friend (or anyone else) that this is the way so- called "perception-behavior-perception loops work (they control perception -- not output) then Bill Powers could have saved 30 or so years and psychology would now be a real science.

John Van Loon (920817) says:

> To me it would seem that drivers do infact see the disturbances.

A good way to see that this is not the case is to look at my demo described in chapter 3 (first paper) in my Mind Readings book. Another way to see this (mentally) is to remember that the visual image through the windshield of the car is ALWAYS simultaneously the result of disturbances AND driver outputs. Even the effect of a sudden, transient disturbance depends to some extent on what the driver (and other disturbances) were doing at the time it occurred. When control is good, responses (outputs) are always the exact opposite of the net effect of disturbances on perception; there is nothing in the perception that could be used to infer the net effect of the disturbance. That is such an amazing fact that even people who understand control theory can hardly believe it (right Gary).

> The action of steering the car so that it stays in the lane
>seems like a negative feedback type of action solely controlled by sight

In fact, this illusion is so compelling that probably fewer than .001% of people currently working in psychology know that it is an illusion. Actions control perception, not vice versa. I know it's hard to believe -- but it's true (and demonstarted over and over again, in many different ways); do read that paper in "Mind Readings"; do the experiment yourself; try to find something in the sensory input that is related to the output that affects that input. You might be startled too -- and see why those of us who understand PCT are so excited about it.

>My

>problem is with the driver not being able to "see" the disturbances and not
>controlling the position of the car on the road. Could you clarify the
>point a little further it would help me believe in your methods and theories

One way to do this is to just look at the equations and believe that they are true (which they are). For example, the closed loop equations say that

o = -kd

Output, o, depends on disturbances, d, (environmental events) -- perception is not part of this relationship (when the loop gain is high enough).

The other way is to set up the demos and run them yourself. Do experiments yourself. Try, as I said, to find the perceptual variable that guides responses in a tracking task -- when control is good. You will find that perception does not guide output in a control loop. Output, however, guides perception. In fact, behavior is the control of perception (is there an echo in here?).

Best regards Rick

Date: Mon Aug 17, 1992 10:55 am PST
Subject: Re: separating action and perception

[From Bill Powers (920817.1130)]

John van Loon (920817) --

Hello, John, welcome to the net as a real live speaker!

>To me it would seem that drivers do in fact see the disturbances. It is
>true that they do not see each individual one but they do see the net
>effect of them by noting the position of the car in the lane (this
>includes any steering errors also).

"Disturbance" is an ambiguous term; it can refer either to the variable that is causing the disturbance (wind velocity, bump in the road) or to the effect that the disturbance has (change in direction of the car's motion). Just lately I'm experimenting with using "independent variable" to mean the CAUSES of disturbances, and "disturbance" to mean the effect, if any.

>The action of steering the car so that it stays in the lane seems like
>a negative feedback type of action solely controlled by sight
>(excepting for large disturbances that change the wheel position and
>the change is "felt" and reacted to).

You're right, of course: this is a negative feedback system, which is what I mean by a control system. The controlled variable, from the point of view of an external observer, is the objective position of the car on the road. But when we look at the situation from the standpoint of the driver, it is a PERCEPTION of the car's position that's being controlled, which is not necessarily the same as the objective position. When the driver sees a satisfactory picture in the windshield, the car may not be centered in its lane. So (and this is the basis thesis of the PCT approach) what the driver is actually controlling is the perception, not its objective counterpart.

Note, by the way, that the perception doesn't control the action; the action controls the perception.

The turning of the steering wheel is based on the difference between the perceptual signal representing the car's position and an internal reference signal, which is more or less a picture of how the road and hood SHOULD look as framed in the windshield (in German, the name for reference signal is Sollwert -- "should-be"). The difference, the error signal, is converted into steering efforts in a direction corresponding to the sign of the error. There are some dynamic filters involved to make the whole loop stable, but that's not necessary for a discussion of the basic idea.

The controlled variable, the car's position, is neither a dependent nor an independent variable. It's part of a causal loop. It depends on the action (applying a torque to the steering wheel) at the same time that the action depends on deviation of the perceived position from its reference state determined by the driver. There are two independent variables in this situation. One is the driver's internal reference signal. The other is the set of all independent variables that can affect the car's position independently of the driver's actions -- what I have been referring to as "the disturbance." I'll say "independent variable" when I refer to the environment, and "reference signal" when I refer to the driver's brain. They're both really independent variables with respect to the steering control system.

What drivers don't see are the independent variables in the environment. They don't see the wind, and even if they see some dust blowing or trees bending, they can't see enough to calculate the forces being generated on the car. They can't see soft tires, little tilts of the road, and so on. All they perceive is where the car is, which they judge relative to where they want to percieve it. If the car swerves, there is no way of knowing what caused the swerve; there might be many simultaneously acting independent variables, which remain inseparable as their effects simply add together, or just one.

This is actually the real power of a control system. It doesn't need to know the causes of disturbances of the controlled variable. It simply monitors the controlled variable itself, compares what it senses with some reference-state for the perception, and acts to keep the difference as small as possible.

The point I was making to Penni Sibun concerned the relationship between the driver's actions (applying torques to the steering wheel) and the perceptions that are controlled as a result (the observed position of the car on the road, as the driver sees it). In some of the materials Penni cited, the statement appears that action and perception are inseparable. But when there are independent variables in the environment which have just as much influence on the outcome as the actions do, it is the outcome that remains the same, while the actions vary to oppose the effects of those independent variables. So the perception (of the car's position) remains essentially stable, but the actions (torques applied to the steering wheel) vary as the independent variables in the environment vary.

As a result, you can be driving down a nice straight road with the steering wheel in just about any position, depending on how much crosswind there is, how much the road is tipped, and how well aligned your front end is. There's no way to tell, just by looking at the car's behavior, what the steering wheel is doing. So action and perception become decoupled; the variations in action correlate mostly with the external independent variables, and hardly at all with the controlled variable. The angle of the wheel correlates highly with the sideward forces due to the crosswind, but because of that, shows very little relationship with the direction the car is going.

The reason that seeing or not seeing independent variables is important is best understood in relation to conventional interpretations of behavior. From the standpoint of an external observer, it seems that the independent variables are causing the behavior -- the crosswind is causing the driver to turn the steering wheel. If you were from Mars, and didn't know anything about driving cars, you might conclude that the driver is somehow sensing the wind, the deviation of the road from level, the state of the tires, and all the other variables than influence the car's direction, and is responding to them by turning the wheel by a calibrated amount. You would see what seems to be a stimulus-response situation.

If you didn't ask too many questions about HOW the driver senses these things, and how he or she does so with such precision, and how these stimuli get translated so precisely into just the steering torques that are needed, you might think you have an adequate explanation of the phenomenon. It would never occur to you that there is another variable involved, the path of the car, that is actually under active control. The fact that the steering forces balance out all the external forces so accurately that the car stays in its lane for mile after mile would not seem remarkable if you weren't a physicist -- that's just good luck for the driver. It wouldn't be at all obvious that what the driver is REALLY sensing is the position of the car; the driver isn't sensing any of those supposed "stimuli."

This, I maintain, is the story of all conventional approaches to understanding behavior. The supposed causes of behavior, the independent variables or "stimuli," are really just influencing perceptions that the organism has under control. The actions that result from applying these independent variables are really opposing disturbances of the controlled variable -- usually very successfully. Organisms are such good control systems, unfortunately, that controlled variables do little varying that can be related to external events. The result is that they're easy to overlook. Statistical analysis is most likely to reject them as insignificant, even if the experimenter accidentally includes them in the list of variables to be tested for significance.

The AI and AL approaches (artificial intelligence and artificial life) uniformly assume that the consequences of acting simply follow regularly from the actions. As in most conventional sciences of life, they name actions by their consequences -- "turning left" rather than "applying a torque to the steering wheel." They would assume, for example, that steering a car around a curve simply results from commanding that the arms turn the steering wheel by a specified amount. As every driver knows, however, it's sometimes both possible and necessary to go around a curve to the left by turning the steering wheel to the right, if the curve is too steeply banked for your speed or if a crosswind is blowing the car into the curve, or both. There's no one "command" that can produce turning left in a real environment. The system has be be organized to turn the wheel by ANY AMOUNT and IN ANY DIRECTION that's required at the moment. Only a control system can behave like that. A command-driven system can't.

I once heard it said that asking Kenneth Orr a question was like trying to get a drink of water from a fire hose. I'm glad I'm not like that.

Best, Bill P.

Date: Mon Aug 17, 1992 11:14 am PST
Subject: Re: independent Investigator #2?

Re: Independent Investigator #2?

From Tom Bourbon (920817)

What an oversight -- Greg Williams has been #2 for years! If I must assume the
title, the company of #s 1 and 2 is tops.

Thanks to everyone (many in private posts) who offered condolences, support, and suggestions about people and places to contact. Keep those cards and letters -- and job prospects -- coming.

Warm regards -- Tom Bourbon Interim (?) Independent Investigator #3

Date: Mon Aug 17, 1992 2:14 pm PST
Subject: Re: PCT vs cognitivism

(penni sibun 920816)

> [marken 920814]
>
> I must also mention that Turvey has the honor of having published the longest, most systematic and most incorrect critique of PCT to date.

i'll believe you. i don't think that detracts from the fact that he's not a behaviorist.

> There are many incarnations of cognitivism. But what is common to all of them (and what is commonly wrong) is that they view behavior as the end result of mental processing. PCT says that it is perception that is the result of "mental events" (reference signals).

i think this charaterization overstates cognitivism. i don't think cognitivism necessarily sees beh. as the ``end result'' of mental processing; i think cognitivism sees mental processing as crucial for
behavior.

> This apparently small
> distinction is as big as the distinction between behaviorism (which sees external events as the cause of behavior) and PCT.

no, it's not: they both have the word ``mental'' in them, and behaviorism clearly does not.  the emphasis on the ``mental'' distinguishes cognitivism and pct from behaviorism.

> There are some cognitive models that have some closed loop characteristics; and probably qualify as control system models of the control of higher level variables. An example is Simon's theor[e]m prover (and current models that are similar, like the prolog language itself). These models try to satisfy certain logical conditions (the reference condition) by trying to satisfy other logical conditions in their database. These are control models that work in a logical (not a real) environment (the database). These are interesting models and they may eventually become useful as models of what Powers calls "program level" control. I have no beef with these models -- I think they are very clever. They are just not yet connected in a realistic way to behavior.

you don't find the fact that they are intractable problematic? interesting (but somewhat off the point).

>>if you think it's too weird to consider that organism and environment
>>are fundamentally inseparable, why not consider a slightly different
>>take on it: action and perception are inseparable; there is no way
>>you can draw a line and say that on one side is the organism's action
>>and on the other is the organism's perception. on this view, what
>>could it mean to locate control in the head of the organism?

> In Bill Powers' "Living control systems" page 251-252, precisely this question is answered, simply and quantitatively. The short answer is "because the organism amplifies, the environment dampens". (my quotes, not Bill's).

this is simply sloganeering; it doesn't address the question in any obvious way (though i realize it's been addressed in later messages).

>>on the interactionist view, control is neither in the head not in the
>>environment.

> Sounds fine; but it turns out to be wrong -- a common problem when trying to understand the world with words instead of models.

do you think you could cut down on the _ad hominem_ rhetoric?

> I would like to learn about interactionism. But I didn't really learn much from the quotes you provided.

fair enough. read almost anything by agre. i can give you preston's address and you can ask for her paper.

> Perhaps you could describe an interactionist model of some very simple phenomenon.

see below about driving.

> Matter of fact, perhaps you might start by explaining the nature of

the phenomenon that interactionism explains.

behavior. the same phenomenon that cogitivism and behaviorism explains.

Does it explain the phenomenon of control? If so, how?

no. it does not seek to explain control. unlike cognitivism and behaviorism, interactionism is not founded on the premise that the idea of control is crucial to understanding or explaining behavior.

[from Avery Andrews (920815)]
(penni sibun 920814, on marken 920808)

What I'd say to this is that PCT is not cognitivist, because the *phenomenon* of control is only found when an organism (or other `Agent') is put in an appropriate environment (e.g. if the road is covered with an oil-slick, one ceases to observe control of the direction in which the car is going). The phenomenon of control is just what Rick says it is: when you introduce disturbances that you would expect to change some aspect of the environment, but that aspect doesn't change, because the Agent consistently does things that `mysteriously' counteract your attempted disturbances.

on an interactionist view, there's nothing mysterious going on. (though, frankly, i must confess that i don't really understand y'all's rhetorical use of ``mysterious''; i don't really know what's supposed to be mysterious to whom.)

PCT goes on to say (at least) two further things:

a) an Agent that is efficacious in the real world must effect control

b) the only way control can be effected is by means of certain kinds of internal arrangments - perceptors, comparators & effectors appropriately connected so as to constitute negative feedback loops in the context of the actual environment.

i think it's conceivable that pct-type control might be more interactionist than any of us understands at this point. however, part b) is squarely cognitivist: it requires an inside-outside line to be drawn, and puts crucial stuff on the inside of the line.

[From Bill Powers (920815.0800)]

Penni Sibun (920814) --

You say some things, and represent others as saying some things, that fall strangely on my ear. For example:

>for these folks
>and others in ai, for gibsonian psychologists, for
>ethnomethodologists, etc., agents and their worlds are *not
>separable*. how can you have a respond to b's stimulus when a and b

>are the same thing?

If they can't separate agents and their worlds, they're trying to use the same model for the inanimate environment and for living systems. This is simply a mistake: the inanimate world does not run by the same rules as the animate one. Of course if you jack yourself up to a high enough level of abstraction, everything is the same as everything.

hold on. doesn't everything bottom out in physics? certainly the current view is that everything is made of the same subatomic building blocks, in a very fundamental, not abstract, way.

But then you have
nothing interesting to say, because all the interesting things we say about ourselves and the world are about relationships between independent things. If you can't tell things apart, because of having abstracted too far, then you can't see any relationships.

ok, i agree that when we say interesting things, we say things about objects and relationships that we have individuated for the purpose of talking about them; certainly there are things that we typically individuate (like organisms) so that we can talk to each other. on the interactionist view, there is nothing sacrosanct about these typical individuations. whenever you look very closely at a clearly drawn line, you begin to see that it gets rather fuzzy. the line between action and perception is a paradigmatic example of this.

It's hard for me to see how
a person who doesn't make this distinction [bet. perception and action] can understand what it means to
say that action controls perception, unless that person means that perception IS action, or that perception follows faithfully from action (which it doesn't -- that's not the meaning of "control").

all true if you assume control. but as i said, the interactionist view doesn't assume a model of control.

But action and perception are completely separable; among our most basic and elementary demonstrations of control processes is one that shows an almost total LACK of correlation between a person's actions and the perception that they control. If you don't understand this fundamental separation, then you've missed a basic point about how control works.

This is such a basic concept that I don't think we can understand each other until we reach some sort of agreement on it.

i understand what you're saying, and i do in fact agree that you have an interesting and efficacious model of control. you can build little creatures who do cool things with your model. you might even be able to build little creatures that are better than anyone else's (though i think the jury's still out on that one).

but i don't think your model is sufficient to account for human behavior (or the behavior of any creature you didn't build yourself).

Think of the car and driver. The driver controls a perception of the car's position in its lane by the action of applying a turning-moment to the steering wheel. Essentially all of the driver's actions are present only to oppose the effects of independent disturbances on the car's path. Those disturbances are unpredictable and invisible. Nevertheless, the driver's actions correlate highly (negatively) with their net vector sum. But the driver's actions correlate hardly at all with the visually-detected position of the car on the road, which is the controlled perception.

How does this jibe with saying that action and perception are inseparable?

i'm not going to answer this question head-on at the moment. your formulation of what's going on is in terms of separate action and perception; it's difficult to answer that question within your framework. let me try to get a step closer by formulating the situation in my own terminology, which is interactionist but is undoubtably somewhat idiosyncratic (there isn't ``standard'' interactionist terminology).

new england is covered in backroads--they are twisty, hilly, narrow, and old, and are subject to nearly all possible weather conditions. the way one drives them is to keep one's left wheels near, or quite possibly over, the center line (if it is visible). essentially, one drives down the middle of the road; this is becuase the danger of running into someone else is much less than sliding ingloriously off the road (when a highschool friend of mine ran off the road shortly after getting his licence, all the adults clucked and opined that obviously he had stayed too close to the edge of the road). the road changes constantly with the weather; if there's been an inch of snow already but nothing's been plowed, driving down the road means driving in the tracks established by previous drivers, even if the tracks are suboptimally placed, since following the tracks is safer than crossing them.

in california, the roads are much newer, and since they'r never plowed, they have little reflecting bumps on the painted road markings. so driving down the road means being mindful of not crossing the bumps (which are quite effective at rattling the car).

the roads in the santa cruz mountains are as twisty as and much steeper than the backroads of new england (but they are better made and have reflector bumps). there are no speed limits in the mountains: if you know the road, you drive fast; if you don't, you drive slow; if you're stupid and drive fast where you don't know the road, you lose. what is crucial is that you stay on yr side of the reflectors no matter what. if you cross the line, any sort of encounter is a disaster. if you stay on your side, the only hazard is rearending bicyclists.

i280 in san mateo county calls itself the world's most beautiful freeway. i don't know about that, but it's certain one of the world's best built, and the traffic's reasonably light. i have the sort of commute that it's easy to zone out on, there's so little to pay attention to. i do tend to notice roadsigns, though, and i know, for instance, that after the woodside exit, the road's going to take a gradual curve to the left, and i can account for that.

ok, that's my car and driver story. i mean it to illustrate that there's no such thing as simply ``someone keeping a car on the road.'' there's always a whole lot more stuff involved. i have mentioned at least the following factors: customs; geography; weather; road construction; individual experience; various sources of information about the road (eg, reflector bumps, road signs). i argue, first, that to model this as a single process of controlling for the perception of a single variable (or a bunch of them) would be hopelessly baroque and probably impossible; second, if you did extract from all this a common thing that we could agree was the essence of ``keeping to the road,'' then you're going to have to tell me where you're going to put all the rest of the stuff that i've mentioned and can produce pretty good evidence for.

the short interactionist story is that none of this stuff gets thrown away, and it's not obviously locatable either Inside or Outside. the long one will have to wait till my hands recover.

Date: Mon Aug 17, 1992 3:43 pm PST
Subject: Re: separating action and perception

(penni sibun 920817)

[From Bill Powers (920817.1130)]

> The point I was making to Penni Sibun concerned the relationship between the driver's actions (applying torques to the steering wheel) and the perceptions that are controlled as a result (the observed position of the car on the road, as the driver sees it). In some of the materials Penni cited, the statement appears that action and perception are inseparable. But when there are independent variables in the environment which have just as much influence on the outcome as the actions do, it is the outcome that remains the same, while the actions vary to oppose the effects of those independent variables. So the perception (of the car's position) remains essentially stable, but the actions (torques applied to the steering wheel) vary as the independent variables in the environment vary.

perhaps you're eliding something here, but you seem to be saying: there's a picture (or encoding thereof) in the head of how the road/hood is supposed to look. perception is the driver getting an unambiguous (and undecomposable?) signal through his eyes, without his having to do any work to get it. there's just tons of literature on perception that suggest that perc. is a *lot* of work; in fact it involves action--perhaps to the degree that p. and a. are inseparable. (``active vision'' is a current buzzphrase.) if i'm driving into the sun, i'll have to squint to see anything. if it's sleeting, i'll have to peer around the bouncing balls of ice and the flailing wipers. at *any* time, my eyes are having to track the road, since the car is constantly bouncing me around. i don't in fact ``see'' everything in front of me the same way; my attention is focused.

it also seems that you assume a one-step comparison process bet. the perception and the reference, but you don't explain how it might work.

cheers. --penni

Date: Mon Aug 17, 1992 3:48 pm PST

From: Dag Forssell / MCI ID: 474-2580
Subject: Dist, Jobs

[From Dag Forssell (920817)]

Bill Powers (920817.1130)

>"Disturbance" is an ambiguous term; it can refer either to the variable
>that is causing the disturbance (wind velocity, bump in the road) or to
>the effect that the disturbance has (change in direction of the car's
>motion). Just lately I'm experimenting with using "independent variable"
>to mean the CAUSES of disturbances, and "disturbance" to mean the
>effect, if any.

Your experiment creates an error signal here. I will grant you that everything is variable, but we have been using the term variable to signify that feature of the environment that we want to control. We talk about the VARIABLE plain and simple. If everything is going to vary, let us come up with a different name.

May I suggest:

```
                      Reference
                          |
                          v
        -------------> Compare ---------------------
        |                                           |
        |                                           v
      Input                                       Output
   (Perception)                                (Instruction)
        ^                                           |
        |                                           v
        |                                      (CS) ACTION
        |                                           |
        |                                           v
        |                                    CS ACTION INFLUENCE
        |                                   (Influence depends on
        |                                 "loose or tight coupling,"
        |                                    -worn linkage, etc
        |
        |              VARIABLE                     |
         -------   Physical equivalent  <-------------------->OTHER
                   to reference signal                  CONSEQUENCES
                           ^
                           |
               DISTURBING ACTION INFLUENCE
                (Influence depends on
                loose or tight coupling)
                           ^
                           |
                   DISTURBING ACTION
                air blast, bulldozer etc
```

The "influences" shown here are only shown as arrows in our graphs, but do often figure prominently in the verbal descriptions. It may not be obvious to all that it is the arrows that are being described sometimes.

Tom Bourbon (920817)

Where does security come from? Is it the art of staying with your ship, (even if it may sink)? Is security knowing how to swim to a new ship?

Tom, have you been swimming any in the last 20+ years? How do you teach Tom to swim? Throw him into the water.

These may not be good times to look for work, but there are some timeless principles to guide you. The principles are to be found in a sometimes misunderstood systems concept called PCT.

Seriosly, looking for work is hard work. There are many books catering to your sense of panic. Some of them are good, but overwhelming in the application. I am thinking of Richard Nelson Bolles' perennial "Parachute," which should be thought of as a reference only (a good one at that).

Presuming that industrial experience can be translated to whatever market you approach, I would heartily recommend:

1) "Executive Jobs Unlimited" by Carl R. Boll. Macmillan 1979.
   Order it, read and see how the suggested activity fishes for open error signals (which you can satisfy) in the people you contact.

2) When you do make a contact, don't blow it by putting yourself down.
   Get "Sweaty Palms; The neglected art of being interviewed" by H. Anthony Medley.

You are embarking on a great marketing adventure. Make the most of it. It can be fun. Knowing that you can do it gives a strong sense of security.

When you post version one, two ..... seven, you will get feedback.

Wishing you all the best,   Dag

Date: Mon Aug 17, 1992 5:31 pm PST
Subject: PCT vs cognitivism

[From Rick Marken (920817.2000)]

I'll try to finish this fast so that I can watch the fascist fandango in Houston -- featuring my favorite actor (really!).

Actually, all I really want to do is straighten out an unfortunate misunderstanding (yes, folks, for a change I have better things to do than post to the net). What I want to straighten out is this:

I said:

>   Sounds fine; but it turns out to be wrong -- a common problem when
>   trying to understand the world with words instead of models.

and penni said:

>do you think you could cut down on the _ad hominem_ rhetoric?

I want to sincerely apologize if this seemed like a personal attack. I did not mean you (penni) personally. I was thinking of theorists (and there are loads of them) who come to conclusions about how things work by "thinking them through" in words, rather than by building working models. This has been a particular problem when it comes to understanding the behavior of closed loop systems (witness the TOTE unit of recent discussion).

I'm glad that you are participating in this conversation; again I apologize if I offended you.

Believe it or not I don't like ad hominem arguments. I did make fun of Turvey -- but that is because he took a rather harsh pot-shot at Powers (in his article) and has never been open to discuss it. I consider that kind of thing crappy -- so I said some mean things about him, but even then, I did it with some reluctance.

CSG net is for discussion and solving problems. Some of the discussions get pretty heated, but I think we have avoided ad hominum attacks pretty well. The reason, I think, is because our best arguments are ultimately solved by demonstration -- showing how an idea works as a working model or experimental demonstration.

Anyway, please accept my apologies for the apparent ad hominum. I don't agree with your position, but I am trying to understand it and to present mine as best I can (within the limits of the medium). Although I think you are wrong, I assume that you are well intentioned. So enjoy the conversation but remember, I can't help thinking that I'm right any more than you can help thinking that you are. We'll both probably be disappointed when we meet Ms. Truth up there.

Best regards Rick

Date: Mon Aug 17, 1992 6:01 pm PST
Subject: Re: PCT vs cognitivism [From Rick Marken (920817.2000)]

s'cool. thanks for reminding me about the show!
--penni

Date: Mon Aug 17, 1992 7:57 pm PST

[from Avery Andrews 920818] (penni sibun 920816)

> i don't really understand y'all's rhetorical use of ``mysterious'';

``mysterious'' means mysterious to a behaviorist or an old-fashioned planning weenie (or, I guess, any random bad guy).

>i think it's conceivable that pct-type control might be more
>interactionist than any of us understands at this point. however,
>part b) is squarely cognitivist: it requires an inside-outside line
>to be drawn, and puts crucial stuff on the inside of the line.

This doesn't fit with my conception of what `cognitivism' is - I would take cognitivism-in-field-X as being the position that everything worth understanding in field X can be understood as a process of building mental representations in the head. This is different from drawing an inside-outside line & locating certain things inside the line. In fact, Chapman and Agre seem to do this: Sonja has a clear inside-outside line, with the visual system & much else located inside it. What matters, I think, is whether you are supposed to be able to understand why the stuff inside does anything useful, without also understanding whats going on outside. This is denied by PCT, and not implied by my (b).

I actually think its essential to find ways of understanding internal structure & relating it to `behavior', since otherwise, because everything after all does bottom out in physics, why don't we all go to the beach and wait to buy the book when Stephen Hawking & Co finish figuring everything out?

Switching topics ...

The driving story illustrates nicely what I at the moment consider to be the biggest potential problem for PCT (ignorance of how perception works is of course another big problem, but it afflicts everybody) - it often seems to happen that one can think of certain abstract & high-level variables as being maintained (car moving down road = 1; car lying on it or beside it as a junk heap = 0; reference level = 1) by means of various relatively concrete and straightforward low level variables as being maintained, but in between there's a vast zone where all sorts of stuff might be going on, which, at least initially, seems to have about as much structure as a plate of spaghetti. e.g. sometimes one is controlling for `car near middle of road', other times `car in snow rut', other times `car to right (or is it left??)' of reflectors, & on top of that there's anticipatory compensation (knowledge of the road), however that works, etc.

A simpler case of this sort of thing is the top-level control of the beerbug: the whole bug's nervous system can be regarded as controlling for the bug having a high energy level (low energy levels trigger `behaviors' that typically wind up having the effect that the energy level gets raised), but the actual way in which the sub-systems in charge of wandering, edge-following, and odor-hunting negotatiate to achieve this end is pretty confusing. I certainly can't be sure that PCT ideas will help in clarifying the workings of this kind of system, though I consider it worth spending some time to try to find out. What my point (b) says is that either these negotations will turn out to be castable in PCT terms, with some benefit derived from this way of looking at them, or the model bug will prove to be too dumb to be viable, and the real bugs that are smart enough to keep themselves alive will have PCT-style internals. Such is the claim, at any rate. Obviously, the jury has barely begun to sit ....

(Bill Powers ???)
(since cognitivists aren't trying to understand behavior, I don't think that

I seem to have mislaid the posting this is a reply to, but as far as I can make out, sonja's limitation to the 8 joystick directions is in no way crucial. She would not require deep modifications if she was supposed to drive, say a hovercraft sled with thruster and rotator engines, able to move in any direction (like the spaceship in Asteroids). The reason these interactive AI gizmos effect control is that what they were designed to to is achieve interesting results under circumstances that change rapidly and unpredictably relative to the amount of time it takes them to do anything significant (e.g. kill a monster, as opposed to move a pixel to the left). The fact that the targets move around unpredictably means that there are unpredictable disturbances in the path from gross output to net result, except that the disturbances that C&A focus on are high-level, distal ones (where the object you're heading for actually is) rather than low level proximal ones (how much torque you get for how much neural current). The same dog, just barking in a different corner of the yard.

As for CGA HiRes, writing with setcolor(BLACK) doesn't effect erasure (in CGA HiRes) , but the setwritemode(1) trick looks like what I was looking for.

I had already drawn the gloomy conclusion that re-writing was the only way to erase in Borland graphics, by looking at the NSCK code and seeing that that seemed to be how Pat & Greg were doing it.

Avery.Andrews@anu.edu.au

Date: Mon Aug 17, 1992 8:30 pm PST

(penni sibun 920817.2100)

[from Avery Andrews 920818]

>i think it's conceivable that pct-type control might be more
>interactionist than any of us understands at this point. however,
>part b) is squarely cognitivist: it requires an inside-outside line
>to be drawn, and puts crucial stuff on the inside of the line.

This doesn't fit with my conception of what `cognitivism' is - I would take cognitivism-in-field-X as being the position that everything worth understanding in field X can be understood as a process of building mental representations in the head. This is different from drawing an inside-outside line & locating certain things inside the line.

but _in extremis_, which is certainly how i'm arguing, just saying something's in the head is drawing a line.

In fact, Chapman and Agre seem to do this: Sonja has a clear inside-outside line, with the visual system & much else located inside it.

well i don't think sonja's implementation really addresses inside/outside (though it does address other things), and i don't think c&a think it does either. i don't think anyone knows how to address this in building frobs right now. sonja's deemphasis on the central system makes it not completely inconsistent w/ interactionism, even if c&a succumbed to linedrawing in building it.

cheers. --penni

Date: Mon Aug 17, 1992 9:42 pm PST
Subject: PCT vs. Cognitivism

[From Bill Powers (920817.2100)]

Penni Sibun (920816) --

Your "interactionist" picture of driving certainly shows how complex and variable it can be. I had written a long and windy response, which I somehow, and fortunately, managed to destroy. I think the most productive step would be just to cite a little bit of your description of driving, and ask you a simple question.

You say

>new england is covered in backroads--they are twisty, hilly, narrow,
>and old, and are subject to nearly all possible weather conditions.
>the way one drives them is to keep one's left wheels near, or quite
>possibly over, the center line (if it is visible).

My question is this: would you explain, in terms of cog. sci, interactionism, or any other approach HOW a driver keeps the "left wheels near, or quite possibly over, the center line (if it is visible)"? By this, I mean just explain how the intention that this result should happen gets turned into the actual happening.

Best Bill P.

Date: Mon Aug 17, 1992 9:59 pm PST

[Avery Andrews 920818:1543]
(penni sibun 920817.2100)

I guess I don't (yet) see the point of not drawing a line between the inside & outside of critters. Maybe Sonja is not the best example of this, because she's not a full-scall critter-in-environment simulation, but Randy Beer's bug is, and there seems to me to be a clear difference between the neural circuits & currents on the inside & the locations of the food-patches, barriers, etc. on the outside. I take Beer's point not to be that there is no inside-outside distinction, but that the explanations for behavioral patterns (at the `molar' level, if I remember my psych. jargon corrrectly) are often to be find in neither place exclusively.

Avery.Andrews@anu.edu.au

Date: Tue Aug 18, 1992 3:19 am PST
Subject: closing the loop

[From: Bruce Nevin (Tue 920818 06:59:09)]

A number of us have wondered if there were some way of closing the loop when we send messages to csg-l. Why doesn't the listserv software include the sender among those to whom it sends the sender's message?

I learned from another listserv group that, in order to be included in the distribution of one's own messages, one can send a message to the listserver (the host running the listserv software) with the following text on the first line:

set csg-l repro

(I have substituted csg-l for the other group's list name, of course. The default is said to be NO REPRO.)

I sent this message to LISTSERV@VMD.CSO.UIUC.EDU (the BitNet address is LISTSERV@UIUCVMD.Bitnet). I received a message confirming that the listserver had accepted the command, and Gary confirms that also. In a following message, I will tell you whether or not it worked--that is, whether or not the listserver send me a copy of this message.

Bruce

Date: Tue Aug 18, 1992 3:34 am PST
Subject: Re: closing the loop

[From: Bruce Nevin (Tue 920818 07:24:34)]

I received a copy of my message before receiving the little notice about it having been distributed. The command works.

Bruce bn@bbn.com

Date: Tue Aug 18, 1992 8:05 am PST
Subject: Re: separating action and perception

[Martin Taylor 920818 11:50]
(Penny Sibun 920817)

> there's just tons of literature on
>perception that suggest that perc. is a *lot* of work; in fact it
>involves action--perhaps to the degree that p. and a. are inseparable.
>(``active vision'' is a current buzzphrase.)

This was cast as a criticism of PCT, but in fact it seems to be a reasonable description of how PCT says perception works. An ECS that is trying to achieve a percept can go only by its error signal. It doesn't know what it is doing when it amplifies that error signal into an output that contributes to the reference signals of unknown lower-level ECSs. All it knows is that the connections are such that producing this output normally leads to a reduction in its error (i.e. a percept more like the percept it is trying to achieve). From here, I can continue your quote to describe the sort of thing that might occur:

" if i'm driving into the sun, i'll have to squint to see anything. if it's sleeting, i'll have to peer around the bouncing balls of ice and the flailing wipers. at *any* time, my eyes are having to track the road, since the car is constantly bouncing me around. "

A good description of the actions of a PCT hierarchy.

Martin

"Statistics are easy to understand. Unfortunately they are even easier to misunderstand." (Me)

Date: Tue Aug 18, 1992 9:16 am PST
Subject: RE:Drivers, cognition, & PCT

[Mary Powers 920818]

(Penni Sibun 920816)

>but I don't think your model is sufficient to account for human behavior
>(or the behavior of any creature you didn't build yourself).

>I argue, first, that to model [the car and driver story] as a single process of >controlling for the perception of a single variable (or a bunch of them) would >be hopelessly baroque and probably impossible; second, if you did extract from >this a common thing we could agree was the essence of "keeping to the road", >then you're going to have to tell me where you're going to put all the rest of >the stuff that I've mentioned and can produce pretty good evidence for.

Can we first take care of cognitivism by saying that mental processes that compute outputs will not get you down that twisty New England road.

Control theory didn't start with building a little creature and then jumping to behavior in general. It started with the question of how humans and others can behave at all. The Little Man (or Arm) simulates the presumed first three levels in a physiologically plausible way (and that probably takes care of the control organization of most of the creatures on this planet).

One thing to remember about control systems is that they are interactive in the sense that the loop is always closed through the environment, which is as much a part of the system as comparators or reference signals. But interactive implies that the road is out there doing something to the driver, and it isn't, it's just lying there.

Modelling your drive as a the action of a control system does seem hopelessly baroque at first glance. All this stuff going on. But this is where the hierarchical model comes in, to organize all that confusion. Customs are programs and sequences (signal for turns, drive on the right, stop at red lights) that satisfy principles (rules of the road) which are the means for keeping most drivers alive. Geography is a highly abstract systems concept derived from perceptions of curves, hills, scenery (no redwoods in New England), things you've read (New England is a broken off hunk of Scotland), etc. Weather is another systems concept - the perception of it's snowing now is part of it, and so is the radio telling you a cold front from Canada is hitting warm air from the South Atlantic on a line from Springfield to Boston. Road construction - as I go I see that everything you name is either a perception like a bumpy road or snow, or a very high level abstraction like geography or weather. No wonder you think that getting from one to the other is beyond the reach of mere control systems. Take "keeping to the road". Ask yourself how? By steering. How? By turning the wheel. How? By moving this arm this way and that arm that way. How? By tensing your muscles. Start again with keeping to the road. Why? To stay out of the ditch? Why? (Aside from not wanting to get killed). Because getting a tow truck etc will make you late for work. Why? You have a lot to do today. Why? Getting it done brings money, praise, and satisfaction. Why? You like to feel good about yourself. Why?

The hows are probably pretty consistent from person to person; they are the means by which you accomplish the whys, but the whys will vary from individual to individual, and in the same person from time to time. The point is not which particular path is followed from muscle tension to self concept, but that the path goes from level to level, as what you are controlling for becomes more and more complex. The higher levels control by setting the reference signals for control systems at the level below, and they set the reference signals for syems at the the level below that, and so on.

Bill suggests eleven levels: intensity, sensation, configuration, transition, event, relationship, category, sequence, program, principle, and system. His little creature incorporates the first three, and the model checks out that far. Several therapists using this approach focus on the last three, and find that it very well accounts for human behavior, misbehavior, and conflicted behavior - the eight (presumed) levels below are the means by which their clients have screwed up their lives - the hows.

I get the feeling that you don't like the idea that people are merely or only control systems. The point is, that is the way they appear to be organized, and it seems to be an organization that is capable of generating all the complexity and richness of human life. Not in one big jump of one great tangle of a control system, but level by level, each feeding into a system that controls a slightly more complex set of variables than the ones below.

[NOTE FROM BILL]

This is one of those "if everybody did it ..] cautions. The repro command sends your file back to you. But if everybody did that, the network capacity would be halved.

Date: Tue Aug 18, 1992 10:19 am PST
Subject: Re: PCT vs cognitivism

[From Rick Marken (920818)]

First some off the wall comments before I try to answer some of penni sibum's specific comments.

1. The fascist fandango was too scarry for me; I had to turn it off after Buchanan's speech. What a sweet, caring, tolerant bunch of what, I suppose, have to be classed as human beings.

2. Did anyone (such as you, Bruce Nevin) invite that editor from Discover to listen in on CSGnet? I think it would be a nice idea to try to get something published in that forum.

penni sibun (920816) says:

> i don't think
> cognitivism necessarily sees beh. as the ``end result'' of mental
> processing; i think cognitivism sees mental processing as crucial for
> behavior.

In the cognitive model I used as an example, a behavior -- a sequence of finger taps -- is thought to be generated by a hierarchical set of commands. These commands are generated by a mental plan (that presumably exists in the brain). As this plan unfolds it generates the appropriate finger sequence (in theory) The movement of each finger produces the sequences of taps which is the behavior being accounted for by the model. So the behavior (sequence of taps) is the end result of the mental plan. It works like this:

mental plan >> efferent commands >> finger movement>> taps

This is a lineal causal process which I argue is characteristic of ALL cognitive models. Note that behavior (in this case, the taps) is at the end of the line.

The problem with this model is that it neglects some important aspects of the situation -- namely, disturbances. Disturbances are known to enter the process right after the efferent commands are generated. Efferent commands are neural impulses that cause muscles to tense; but the amount of tension per unit efferent signal varies considerably over time; so the result of efferent commands is disturbed by variations in the effect of neural commands on muscles. Environmental disturbances enter the picture after muscle tensions occur and influence the effect of these tensions on the results they are intended to produce. For example, changes in the initial position of the finger will influence whether or not a particular amout of tension will result in a tap or not.

The existence of disturbances means that there is no way for a system to reliably produce an intended result via the same means (in this case, the same efferent commands). This meams that there is no question that Rosenbaum's sequence production model will not behave as the subjects do in a real environment (the same environment as the one the subjects work in).So this type of cognitive model is unquestionably wrong.

The only way that a system can be designed to produce reliable results in a disturbance prone enviroment is to design the system so that is controls the perceptual representation of those results; this sort of model is based on a closed loop organization -- it is a control model. The important characteristicof such a model is that is varies its output (efferent commands) in order to keep the perceptual representation of the result at an internally specified reference level; it controls its perception. Once you understand this, then you realize that the process of understanding behavior is not a matter of learning how systems generate complex sequences of outputs. It is a matter of learning what aspects of their own perceptual experience are being controlled. Behavior is organized around the control of perception; indeed, I would say that behavior IS perception in action (a slogan, I know, but one that makes sense when you watch the behavior of a control hierarchy, especially from the point of view of the control hierarchy itself). If you think about it, you will realize that your own behavior, from moving your fingers to trying to argue the virtues of interactionism, is a matter of producing the perceptions that you want to produce. In fact, there is no behavior without perception; I argue that you can't do anything (intent-ionally) that you cannot perceive (I know that we do a lot without attending; but how much do you do without perceiving?). I know that there are some cases where behavior seems to happen without perception; for example, you can move your leg even though you cannot perceive it because you've slept on it all night. But I would argue that what you are controlling is perceptual side effects produced by throwing the leg around basically; you certainly can't control the degree of bend real well, if you can't perceive it.

The fact that perception is fundemental to PCT means that people who have always been students of perception (like myself) have a whole new level of justification for their interest. The study of perception is no longer just an arcane little sidelight to understanding behavior; it's what behavior is all about. So it is strange to me to hear you say:

penni sibun (920817)

>there's just tons of literature on
>perception that suggest that perc. is a *lot* of work; in fact it
> involves action--perhaps to the degree that p. and a. are inseparable.
(``active vision'' is a current buzzphrase.)

All the literature in the world won't convince me that I'm working hard as I glance around and perceive different aspects of the world (my neural networks may be computing some very complex functions of my sensory inputs, but it feels pretty effortless). There's different levels of bright- ness, lots of colors and sounds, patterns, objects (here's a computer screen, there's a bunch of books), relationships (I'm closer to the computer than to the door), events (the secretary just when to the other office), sequences. I can even perceive principles and system concepts with little effort (I was perceiving some pretty ugly ones last night, that's for sure -- though, of course, they were only ugly relative to my own references; the people at the convention seemed to think they were perceiving great principles).

While it takes no effort to perceive, it can take substantial effort to con- trol those perceptions (get them to the level you want). If I want to perceive that big book over here rather than over there I'm going to have to exert some energy to change that perception. Controlling principles and system concpets can require some real heavy duty effort; but perceiving the principle or system concept seem seffortless to me.

I have to imagine that active vision is controlled visual variables-- and it seems to be from your examples (like squinting to reduce glare). Maybe these active vision people are, indeed, doing PCT.

Anyway, back to controlling other perceptions.

Regards Rick

Date: Tue Aug 18, 1992 10:42 am PST
Subject: Re: ???

I am not a diehard PCT guy (have not set to really exploring the possibilities, yet) but I can see the imediate advantages in control theory. Recently I looked over the ??? notes and saw a possible connection to other ideas circulating arround.

from Avery Andrews 920818

>The driving story illustrates nicely what I at the moment consider to be the
>biggest potential problem for PCT (ignorance of how perception works is of
>course another big problem, but it afflicts everybody) - it often seems to
>happen that one can think of certain abstract & high-level variables as being

>maintained (car moving down road = 1; car lying on it or beside it as a junk >heap = 0; reference level = 1) by means of various relatively concrete and >straightforward low level variables as being maintained, but in between there's >a vast zone where all sorts of stuff might be going on, which, at least >initially, seems to have about as much structure as a plate of spaghetti. e.g. >sometimes one is controlling for `car near middle of road', other times `car in >snow rut', other times `car to right (or is it left??)' of reflectors, & on top >of that there's anticipatory compensation (knowledge of the road), however that >works, etc.

The reference signal does not have to be "digital". I personally would set it to an analog function. Instead of the car being on or off the road I prefer to say where on the road "I" want the car. Ie: -1 for the ditch on the left and 1 for the ditch on the right (-2 ... offroading?) with values in between.

The problem now exists of how to determine what the "best" refrence signal or even a good enough refrence signal (position on the road/ditch) is. This, I think, could fall into pattern recognition which can be handled by neural networks (tedious and difficult ... yes! but possible). The networks can be taught to generalize as to the conditions/customs of driving. If the "car" was transplanted into a new enviornment it may crash or it may "learn" (just like a tourist), the only way to tell is to try it out (in simulation or at worst, not in my neighbourhood).

Getting away from networks, if more sensors and input devices are put in place to deal with narrow roads, high speeds, memory input, etc then these can be incorporated into the feedback loop by changing the reference signal, or adding to the disturbance signal. This will mean that instead of trying to drive down the right side of the road when in england memory input changes the reference signal (*[-1]) so that the car goes down the left side, or the middle, or arround some oil or other substance (would you really avoid oil if you saw some liquid on the road?). If one control loop is responsible for the control of the car, can`t another be in control of the first? I am getting too far ahead of myself, I hope that I have made some points that may have been unclear.

I would like to thank B. Powers for that vapourous explanation kindly given and wish Tom good luck (that is what I get to look forward to in 8 months). It was very helpful, now it is time to pick up some more lititure and try to digest all that I have swallowed here (not water) in the last couple weeks. This term I will be building an autonomous vehicle. I will try to organize it to accept PCT concepts and hopefully it will serve as a viable testing ground for different theorys. Should anyone have suggestions pertaining to traps to avoid or areas that need to be addressed My ears are always open and any critisim (or that elusive praise) is more than welcome. Incidently the platform base will probably be one of those battery operated cars from Toys Are Us (not my idea or first choice but it will save time, maybe). A great way to test out some of these "Driving Miss Daisy" ideas.

-John -a tenderfoot in the way of electron behavior

John van Loon: John_Van_Loon.XRCC%xerox.com@uunet.ca

Date: Tue Aug 18, 1992 11:00 am PST
Subject: Re:re: ???

Sorry to jump the gun on my last posting. I had forgotten to check my mail after lunch and missed RE:Drivers, cognition, & PCT posting. On the subject of the Repo command, how can I get that thing turned off? It is most annoying.

Date: Tue Aug 18, 1992 11:16 am PST
Subject: interactionism

(penni sibun 920818.1200)

i'll reiterate two things up front:

1) i'm not defending cognitivism; interactionism isn't cognitivism (or behaviorism!).

2) i'm quite willing to believe that as far as a technology for building things goes, pct might be the best candidate; i do like the feel of the model in a lot of respects.

> [From Bill Powers (920817.2100)]
>
> My question is this: would you explain, in terms of cog. sci, interactionism, or any other approach HOW a driver keeps the "left wheels near, or quite possibly over, the center line (if it is visible)"? By this, I mean just explain how the intention that this result should happen gets turned into the actual happening.

no i can't explain how; i couldn't build or design something to do it. i can say that i don't think it's necessarily the right question. in other words, if we could account for all the things that i've mentioned (customary ways of driving, road construction, etc.), then maybe we don't even need to posit an *intention* to keep on the road. maybe, given everything else, keeping on the road is just the easiest thing to do. (i'm not arguing that organisms always do the easiest thing.) that's what interactionist approaches try to get at: an organism doesn't have to posit its goals _de novo_ and figure out how the satisfy them. there's already a lot of stuff around that facilitate doing what needs to get done. as agre puts it, we ``lean on the world'': our roads and cars are culture are designed to make driving down the road a plausibly easy thing to do.

> [Avery Andrews 920818:1543]
>
> I guess I don't (yet) see the point of not drawing a line between the inside & outside of critters. [...]
> the food-patches, barriers, etc. on the outside. I take Beer's point not to be that there is no inside-outside distinction, but that the explanations for behavioral patterns (at the `molar' level, if I remember my psych. jargon corrrectly) are often to be find in neither place exclusively.

i agree. the rhetorical point of saying that there aren't any ``real'' lines is to make us think when we draw them, and be open to the possibility that it might be sensible to draw them in a different place for a diff explanation or a diff model.

[Martin Taylor 920818 11:50]
(Penn[i] Sibun 920817)

> there's just tons of literature on
>perception that suggest that perc. is a *lot* of work; in fact it
>involves action--perhaps to the degree that p. and a. are inseparable.
>(``active vision'' is a current buzzphrase.)

This was cast as a criticism of PCT, but in fact it seems to be a reasonable
description of how PCT says perception works. An ECS that is trying to achieve a percept can go only by its error signal.

maybe someone can explain this to me. when you say ``signal'' or ``variable'' or ``percept'', it has connotations to me of a unified thing, like a tone, or a light intensity. but when i look at the road, i am not perceiving something like a tone. if you can explain how all the stuff my eyes take in can be a single unified thing, maybe i won't find it so oversimplified.

cheers. --penni

Date: Tue Aug 18, 1992 11:26 am PST
Subject: impact of repro command

[From: Bruce Nevin (Tue 920818 14:38:05)]

Bill says:

>The repro command
>sends your file back to you. But if everybody did that, the network
>capacity would be halved.

You'll have to explain your reasoning.

As I see it, CSG-L is not a network. It is a set of subscribers on a number of networks in or connected to the Internet (BitNet, MilNet, CSnet, ARPAnet [is that still what the .com domain is called?], and so on). The *capacity* of each network to bear traffic is not affected by the listserver sending one more copy of each message that it distributes. Nor is the capacity of the networks, nor of the listserv software nor the host that it is running on appreciably challenged by this increment of email traffic--if it were, we should worry right away about limiting the growth of CSG-L, and send only short messages at carefully coordinated intervals. Before that, many system administrators at many sites would be budgeting to purchase increased computing capacity, and network analysts would be urging the purchase and installation of more network nodes to alleviate bottlenecks.

Am I missing your point here? Bruce

Date: Tue Aug 18, 1992 11:59 am PST
From: PAPANICOLAOU
EMS: INTERNET / MCI ID: 376-5414
MBX: PAPANICOLAOU@beach.utmb.edu

TO: * Dag Forssell / MCI ID: 474-2580
Subject: I. I. #3

To: Dag Forssell <0004742580@MCIMAIL.COM>
From: Tom Bourbon
18 August 1992

Dag,

Thanks for your post and for the suggestions.

I was not really in a state of panic, but more a state of ignorance about what, if anything, might be available. For over 25 years, I was tucked away at an obscure university in a small town in the deep woods of East Texas. Consequently, I had no feel for the market -- as you see, my first venture away from academe was not very well informed!

People from CSG-L have suggested a surprising array of leads that I am following up; but I must admit that the most appealing route is the least secure of all. I am tempted to work with Greg on starting the journal for CSG, and to have a try at setting up courses or workshops on PCT, modeled after the summer course in Aix, but on a much less grandiose scale.

Best wishes, Tom Bourbon

Date: Tue Aug 18, 1992 12:13 pm PST
Subject: Re: impact of repro command

[from Gary Cziko 920818.1453]

Bruce Nevin (Tue 920818 14:38:05) says:

>
>Bill says:
>
>>The repro command
>>sends your file back to you. But if everybody did that, the network
>>capacity would be halved.
>
>You'll have to explain your reasoning.

Doesn't make much sense to me either. If there are 101 people on the CSGnet, if person 1 sends a post, it goes out to 100 other people (default is not to send post back to sender). If person 1 uses the repro option to receive his own post back, the post now goes out to 101 people instead of 100. This is a 1% increase in traffic, not a 100% increase. If everyone on the list uses the repro option, the result is the same, 101 message instead of 100 for each post.

Moral: Use the repro option of you want your own messages back. But then to be consistent you should also be leaving messages for yourself on your answering machine

Just another reason to be suspicious of the mathematical reasoning of physicists! --Gary

Date: Tue Aug 18, 1992 12:49 pm PST
Subject: interactionism

[From Rick Marken (920818.1300)]

penni sibun (920818.1200) says:

>that's what interactionist approaches try to get at: an
>organism doesn't have to posit its goals _de novo_ and figure out how
>the satisfy them. there's already a lot of stuff around that
>facilitate doing what needs to get done. as agre puts it, we ``lean
>on the world'': our roads and cars are culture are designed to make
>driving down the road a plausibly easy thing to do.

I think the we (PCTers) and the interactionists may just have a fundementally different notion of what it means to understand something. I don't get any feeling of greater understanding about how behavior works from the suggestion that there is already a lot of stuff around that facilitates doing what needs to be done. The way I parse it, it sounds like you are saying that the environment (or our perception thereof) -- which is the stuff around-- guides (facilitates) behavior (what needs to be done). This sounds like a verbal version of the sr explanation of behavior that you clain interactionism is not. Saying that we "lean on the world" doesn't help; suppose I want to (need to) point straight forward? I can't lean on gravity to do that; if I lean on gravity I'll generally end up pointing down. I really just don't get this interacting business. As Mary Powers pointed out, interacting suggests that the actor and environment are cooperating to produce what needs to get done. But the enviroment doesn't care whether what needs to get done, gets done or not; the environment is just there, doing its own thing (I think; at least that's the model of the environment that we get from physics -- and it works extremely well). The road doesn't care whether you stay on it or not. Whether the road makes it easy to drive or not depends on what YOU need to get done; some of these nice, smooth roads may make it hard to give the kids a fun bump in the back seat. Whether the environment is easy or not makes sense only inb terms of one's aims in that environment; I think.

>maybe someone can explain this to me. when you say ``signal'' or
>``variable'' or ``percept'', it has connotations to me of a unified
>thing, like a tone, or a light intensity. but when i look at the
>road, i am not perceiving something like a tone. if you can explain
>how all the stuff my eyes take in can be a single unified thing, maybe
>i won't find it so oversimplified.

I think this is an important point. I think a number of people have a problem with this. I take it for granted that perception is just what afferent neural firing looks like when you are a brain (which we are). I imagine that every different perception we have is the firing rate of a single neuron. My mental model of this is the receptive field. We know that certain neurons (in the lateral geniculate nucleus, for example) are "looking" for particular patterns of light on the retina. For example, the firing rate of a particular neuron might increase as a line on the retina is rotated from vertical to horizontal. So this neuron is a horizontal line detector. The faster this neuron fires, the more the perception it is having is like a horizontal line. The orientation of the line on the retina is the input variable, the rate of firing of the neuron is the perceptual variable. Subjectively, I imagine that the change in firing rate is experienced as a change in orientation. The rate of neural firing, by the way, is the "perceptual signal" and we also call it the perceptual variable (because it varies). Obviously, I could control this variable if I could influence the orientation on the retina of the cause of the firing in this neuron; and this is, indeed, how neural control systems work. As I rotate my head (or an obhect in the world) I can bring the firing rate of the perceptual neuron (the perecptual signal) to the reference level that I specify. Note that the reference level is also a neural firing rate; all I'm doing is telling (with the reference rate) at which rate the perceptual neuron should fire. The consequence of bringing the perecptual neural rate to the reference heural rate is to create some objective state of affairs -- such as orienting the horizontal of the computer screen with the horizontal line connecting my eyes.

I imagine that there are many (millions?) of neurons that detect all kinds of different properties of the world -- simultaneously. How you wire up a neuron to sensors so that it fires in proportion to, say, the degree of honesty in a relationship, we don't know. But we do imagine that something like this must be what is done by our own brain. I imagine that when I perceive a person as honest, that very abstract, temporally and spatially defined percept is computed by the neural nets in my brain and results in a level of firing in some neuron (the honesty detector?); and that level of firing is what I experience as the perception of honesty.

This is the part of the model that gets complicated (and interesting) -- for a nicer development see Powers "Behavior: The control of perception".

Best regards Rick

Date: Tue Aug 18, 1992 1:04 pm PST
Subject: constructors/ not detectors

[From Rick Marken (920818.1400)]

Epistimological warning--

When I called the neuron a "detector" I implied that neurons are designed to detect things that are actually "out there". This is an unnecessary assumption. Neuron's can be hooked up to sensors in such a way that they computer ANY kind of function of the sensory input. So it's wrong to call them detectors -- better to call them "constructors" because they are constructing perceptions from the raw material of the senses.

This fact becomes important in PCT 201. So we might as well get it right in the introductory course.

Best regards Rick

Date: Tue Aug 18, 1992 1:42 pm PST
Subject: Re: PCT vs cognitivism

[from Gary Cziko 920818.1632]

Rick Marken (920818) asks:

>2. Did anyone (such as you, Bruce Nevin) invite that editor from Discover
>to listen in on CSGnet? I think it would be a nice idea to try to get
>something published in that forum.

Thanks for the reminder. I just sent Carl Zimmer some introductory propaganda on PCT and CSGnet. We'll see if he bites.--Gary

Date: Tue Aug 18, 1992 6:37 pm PST

[Avery Andrews 9108191231]

(penni sibun 920818.1200)

>maybe someone can explain this to me. when you say ``signal'' or
>``variable'' or ``percept'', it has connotations to me of a unified
>thing, like a tone, or a light intensity. but when i look at the
>road, i am not perceiving something like a tone. if you can explain
>how all the stuff my eyes take in can be a single unified thing, maybe
>i won't find it so oversimplified.

With a sonja-like visual system, one might set one marker to tracking the hood-ornament, another the center line of the road, & let the controlled perception be the distance between them (or so I guess--having never built a sonja-like visual system, I don't have full faith in my intuitions about them).

On the general subject of the supportiveness of the environment, I suspect that the difference between PCT & c&a-style interactionism might be more one of rhetorical emphasis than substance. PCT in fact depends very much on the fact that many (in fact most, by an overwhelmingly large margin) features of the environment can be leaned on. A steering subsystem would for example depend on all sorts of facts about roads and carparts about which it knows nothing, as well as upon the workings of the lower-level control-systems it works through. To perform at a minimal level, all that it has to `know' is the relationship between change in torque-applied-to-the-wheel to change in relationship- between- hood- ornament- and- centerline- markers (of course, for hi performance, lots more is needed, but the extra knowledge isn't about car parts, etc.).

One theme of c&a interactionism, as I hear it so far, is that you don't have to know very much to get by -- keeping your eye on the situation and following a few simple rules is enough. PCT claims that you can say a bit (or maybe a lot) more than that, in particular, something about the general nature of the kinds of rules that will in fact suffice to get you by -- that on the whole, they tend to be such as to keep perceptions (= the output of `registrars') at particular values, or at least within particular ranges of values. The claim is that architectures that don't do this won't work. I would take this claim to be decisively refuted if someone built a robot vehicle that did stay on the road in an interactionist manner, but just because it was `the easiest thing to do', without containing anything remotely like a control system controlling a perception along the lines of `position-of-vehicle-relative-to-the-road'.

It remains to be seen if this is a `golden thread' that will be the indispensible key for untangling or building these systems, or if it is just one of a dozen rules of thumb for doing so. It does so far seem to me to be a useful organizing principle that is absent from c&a (but for all I know, this could be only because they think it's so obvious that it isn't worth mentioning). And it also remains to be seen if, when applied complex systems, it will apply cleanly, or instead grow so much hair that it will get difficult to discern the original beat under the fuzz.

For example, whatabout the little positive feedback loop in the beerbug's feeding controller (based on the supposed wiring of the infamous Aplysia). What goes on here is that if the bug is hungry enough (the energy level is far off enough from its reference level) and over food, it will start eating, but being eating also excites the feeding controller, so it keeps eating even when its energy level is no longer low enough to successfully trigger eating. The whole system can be regarded as (part of) a control system for the maintenance of energy-level, but the positive feedback loop looks like at odd bit stuck in.

Avery.Andrews@anu.edu.au

Date: Tue Aug 18, 1992 7:41 pm PST
Subject: Drivers, cognition, & PCT and perception

(penni sibun 920818.2000)

> [Mary Powers 920818]
>
> But interactive implies that
> the road is out there doing something to the driver, and it
> isn't, it's just lying there.

the road may not be *doing* anything, but it's constantly changing, and it's an artifact--somebody designed it and built it--and these are important thing to know about it.

> Modelling your drive as a the action of a control system does
> seem hopelessly baroque at first glance. All this stuff going on.
> But this is where the hierarchical model comes in, to organize
> all that confusion.

i am deeply suspicious of hierarchical organization from my experience in modelling language (as well as other behavior). language is supposed to have all sorts of hierarchy in it (eg, a text is composed of paragraphs are composed of sentences are composed of words) and i just don't think this is a very useful way of looking at language, and am working to show this.

> Customs are programs and sequences (signal
> for turns, drive on the right, stop at red lights) that satisfy
> principles (rules of the road) which are the means for keeping
> most drivers alive.

i really can't imagine how a custom can be a program (in somebody's head). assuming it could be decomposed into a program, how does an individual introspect about it? how are customs passed on and maintained? how are they recorded? how are they incorporated into things like building (better) roads?

> Road construction - as I go I see that everything you
> name is either a perception like a bumpy road or snow, or a very
> high level abstraction like geography or weather.

you and i are talking about weather in the abstract (and i believe i included it in the enumeration at the end of my story as a summary). when i am driving, snow on the road or rain on the windscreen is extremely immediate and concrete. i am not dealing w/ ``weather'' when i'm driving, i'm dealing w/ this stuff that's in my face.

> I get the feeling that you don't like the idea that people are
> merely or only control systems.

no. i don't have any particular emotional investment in what people are made of. i was the one who suggested it's all physics at the bottom.

i am, as we all are who are engaged in this conversation, very concerned about trying to ask and answer the right questions.

> [From Rick Marken (920818)]
>
> >there's just tons of literature on
> >perception that suggest that perc. is a *lot* of work; in fact it
> > involves action--perhaps to the degree that p. and a. are inseparable.
> (``active vision'' is a current buzzphrase.)
>
> All the literature in the world won't convince me that I'm working hard as
I
> glance around and perceive different aspects of the world (my neural
> networks may be computing some very complex functions of my sensory
> inputs, but it feels pretty effortless).

i am quite confounded to hear this. you really think that how much work a process takes is determinable by introspection? that doesn't sound very scientific to me.

> I have to imagine that active vision is controlled visual variables--
> and it seems to be from your examples (like squinting to reduce glare).
> Maybe these active vision people are, indeed, doing PCT.

i was using very gross examples, where i thought the work involved would be obvious, that is, it involves muscles, rather than primarily neurons. active vision folks actually focus more on where the work involved in perception is not as obvious as muscle movements (though it can light up half your brain on the scanners!). the paradigm example of this is object perception. going from the retinal map to seeing objects is not a higher cognitive function--the process of perceiving, even in the eye, never mind the brain, is extremely complex, and not at all a matter of passing an uninterpreted signal down the line. at least, so say the tons of literature, and our resident psychologist, whom i consulted just in case i had missed a major shift in the field.

[From Rick Marken (920818.1300)]

around that facilitates doing what needs to be done. The way I parse it, it sounds like you are saying that the environment (or our perception thereof) -- which is the stuff around-- guides (facilitates) behavior (what needs to be done).

not really. remember, i started out denying that there was such a thing as an ``environment.'' and i also denied that there is a locus of control anywhere.

as an aside, i don't think behavior is ``what *needs* to be done.'' i think it's what *is* done.

This sounds like a verbal version of the sr
explanation of behavior that you clain interactionism is not. Sayingthat we "lean on the world" doesn't help;

we lean on the world, we lean on each other, the world leans on us--everything is incredibly well-connected and articulated.

suppose I want to (need to) point
straight forward? I can't lean on gravity to do that; if I lean on gravity I'll generally end up pointing down.

what? ``leaning on the world'' is a metaphor; that's why it was in quotes.

As Mary Powers pointed out, interacting suggests
that the actor and environment are cooperating to produce what needs to get done.

yup. cooperation doesn't imply intentionality. (or if it does, substitute ``work together'' or something.)

But the enviroment doesn't care whether what needs to get
done, gets done or not; the environment is just there, doing its own thing

i'm not sure caring is germane, but interactionism doesn't have a concept of an ``environment'' that is an inert, alien (y'all keep say it's governed by physics and organisms aren't), thing. at it also isn't particularly concerned w/ intentionality--which you seem to be using to mean something very close to control.

I think this is an important point. I think a number of people have a

problem with this. I take it for granted that perception is just what
afferent neural firing looks like when you are a brain (which we
are).

i am *not* a brain. and though i've never seen you, you certainly seem to be possessed of all the fingers and toes, as it were.

cheers. --penni

Date: Tue Aug 18, 1992 7:47 pm PST
Subject: Re: cognitivism

[Martin Taylor 920818 18:30]

CSG-L readers might be amused by this paragraph from a mailing on the Usenet newsgroup "sci.cognitive". I append the author's address in case someone want to point him here (I don't know the usenet name for the extension of CSG-L). Note that this is just one paragraph out of a longish posting. I don't know what HDP (hierarchic dynamic programming) entails, but from a skim of a couple of earlier messages in this thread, it seems somewhat similar to the higher levels of PCT. Could be wrong on that. Here's the paragraph.
==========

Dr. Thompson also comments that he doesn't see a role for "goals" in the AI sense in biological systems. In order to have a clear goals in HDP, you need to be informed when you have reached certain terminal states. I agree that biological systems do not have those terminal states, however, that doesn't prevent them from approximating such states. I suspect there is a calibration process involved.

--
Harry Erwin
Internet: erwin@trwacs.fp.trw.com

=============
Rick...Don't let the loose canon explode the wrong way, please!

Martin

Date: Tue Aug 18, 1992 7:49 pm PST
Subject: Hierarchic Dynamic Programming

[Martin Taylor 920818 18:40]

Just after posting that quoted paragraph from sci.cognitive, I looked at the next posting. I don't think HDP is like PCT at all, but judge for yourselves. Here's the posting:

Martin
================

Heuristic Dynamic Programming in a Realistic Biological Context
Harry R. Erwin

erwin@trwacs.fp.trw.com

As I showed at the 1982 Animal Behavior Workshop in Guelph, Ontario, the optimum strategy for playing a discrete game against nature involving information collection is a simple threshold strategy. The player uses Bayesian statistics to maintain an estimate of his probability of success, and compares that estimate against a threshold at each decision point. If the probability of success remains above the threshold, he continues the game; otherwise, he quits. The threshold can be calculated by treating the game as a problem in dynamic programming. (John Bather, Pers. Com., 1983)

In a biological context, this strategy lends itself to implemention using HDP. The critic network would provide the current threshold value as a local goal value, and the action network would compare the current probability against that value. If the current probability exceeded the threshold, the preferred action would be to continue to collect information; otherwise it would be to quit. Note that the critic network responds to the perceived payoffs and risks of the game and not to the current situation. Both critic and action networks would be prior to the motor cortex, which would then treat both as a combined critic network and attempt to reduce fear to nominal levels.

```
Current payoffs---\
                   O-- local goal value--------------\
Target category---/ (A)               feedback        \
                                         ---           \
Target condition--\                     |   |           \
                   \                    V   |            \ (D)
Self condition----->0--initial estimate --->0-current est>0
                   / (B)                   / (C)           \
Environment-------/                       /                 \
                                         /                decision
Information collected and processed-----/         (expressed as fear level)
                                                               \ (E)
Motor options-------------------------------------------------->0->motor
                                                                   cortex
```

Note that there are a number of places where training would occur. Subsystem A needs to learn how to calculate the local goal values corresponding to various payoffs and intensities of the game (primarily defined by target category). I suspect most species have this hard-coded in the genome. (The local goal values are not obvious functions of the inputs!) Subsystem B can be trained more easily--in mammals, that is part of the role of play and parental teaching. Subsystems C and D are probably hard-coded, even in man. Subsystem C implements logistic functions, while Subsystem D does a simple comparison. Subsystem E probably uses fear level to affect the preference functions for various actions used by the motor controller, although it may select a desired fear level and output partials to the motor controller instead. (I suspect that version is more correct, because the corresponding 2-person game can't be handled by outputting simple fear level, and man does play the 2-person game.)

Cheers, Harry Erwin
Internet: erwin@trwacs.fp.trw.com

Date: Tue Aug 18, 1992 7:49 pm PST

Subject: Introducing Jeff

Introducing Jeff Hunter.

Hello CGS folks. I'm a contractor working within earshot of Martin. Naturally I have been evangelized about (and converted to) PCT.

My PCT credentials:

I've been a lurker on the CSG net for many months. I've also read Powers' Byte articles, but not any of his books. Martin (Taylor) , Allan (Randall), and I have hashed over PCT a fair bit over coffee. I've also looked over the shoulder of Chris Love while he has built the Little Baby (the learning Little Man). I hand-coded a set of ECS's (Elementary Control Systems) that let the Baby put finger to target (purely for testing the software, not to cheat on the learning side).

My PCT motivation:

I am most interested in PCT as a method of partitioning the learning of medium sized nets (100 nodes) on non-trivial tasks. To explain: the size of neural nets is hitting a (soft) ceiling. If you double the number of nodes (to handle a more complex task) you more than double the training time. The current wisdom on solving this problem is to divide the net in half, train each half of the net on half the task, and then add some more net "on top" that takes two half-answers and makes one whole. Naturally this only works in some domains, but it seems to be a win.

The big drawback is that the designer of the net must cut the problem domain in half "by hand".

PCT seems to be a natural fit for this technique. We start with some low-level control systems and train them on simple tasks. Then we add a higher level control system and train it to use the lower levels on a more complicated task. Repeat ad nauseum.

If we can get re-organization figured out we may not even have to partition the task by hand. (This is a bit optimistic. When teaching human children we usually start them simple and work up.)

My projected PCT postings:

I have a few new contributions (of varying profundity) to the ongoing debate. I also have a lot of backlogged followups to postings.

I've done a few quick-and-dirty programs in my spare time. I may get to spend a while concentrating on PCT models at work.

... Jeff

De apibus semper dubitandum est - Winni Ille Pu

Date: Tue Aug 18, 1992 7:52 pm PST
Subject: top-level

Here's a new topic, related to re-organisation.

I started wondering recently how a top-level Elementary Control System (ECS) can remain connected to reality.

To explain: let us take a high-level ECS in a Control Net. (A high-level node is many levels from the raw input and output.) Assume the net starts out untrained (or only partially trained) for its environment. Finally we assume that random re-organization is a major part of its training.

Now the high-level ECS doesn't know what its inputs or its reference mean. All it must do is control that they match. It may be initially set up with input of (target-position - finger-position) and reference of (0). However after a few random re-organizations the input weight for "target-position" may have been set to zero, and the input weight from "elbow-angle" to a positive value. This leaves the ECS training to control "elbow-angle + finger-position" = 0. There is no way for the ECS (or for the random re-organisation) to know that this new function is nonsense.

In general it seems impossible to keep the input "relevant" to the reference without forcing it in some fashion (and thus adding another set of properties to the ECS).

One approach to "forcing it" is found in our Little Baby (a learning version of the Little Man). As in the Little Man the high-level references involve the distance of the finger from the target (as perceived in the right and left retinas).

The Baby has one (or more) layers of ECSs attached to the outputs of its high-level ECSs. However the inputs are connected directly to the Baby's inputs.

```
            R*      <- top-level reference
                   |
                --------
                | ECSs |   <- top-level ECS
                --------
                /     \
               |       --------
      direct ->|       | ECSs |   <-untrained ECSs
      input    |       --------
      to top   |       /     \
               |      /       \
      ...................................
               \   /           \
                Input          Output     <- environment
```

The Little Baby is forced to learn to follow the target by being provided with a fixed input function. The Complex Environmental Variable (CEV) that the Baby is controlling cannot be unlearned, however it can likewise never be learned. This is a reasonable hack while we experiment with re-organisation, but in the long run we can't always hand-code/hard-code the inputs.

Bill seems to have also seen the problem since he has suggested that the learning mechanism should not be completely blind. He wants it to contain some simple CEVs (an oxymoron :?) which guide the re-organisation. This may be necessary, but it also feels like a hack to have a separate control hierarchy for learning.

I have a partial solution that does not add new variables or structure to the existing hierarchy. Unfortunately the CEV in the example is different than "finger on target".

Suppose we have Little Baby (Mark MCXLI) that can successfully learn to control (i.e. we have solved some of the re-organisation problem).

Now we wish to teach it to avoid a spot in its environment (say the exact center of its cube).

We add an extra input (called Pain). We change the environment so that Pain becomes large if the finger is close to the center of the cube, but is very small elsewhere. (We now have a hot-spot.)

We also add a simple ECS that has Pain as input, zero as reference, a large gain, and outputs to the arm muscles. This does nothing while the finger is outside the hot-spot. If the Baby moves the finger into the hot-spot this ECS will quickly yank it out, and will then resume doing nothing. We have given Little Baby a pain reflex.

The Baby now avoids the hotspot very effectively, however it will have trouble moving finger to target in some cases (assume for the moment that we don't move the target into the hotspot). If a trajectory goes through the hotspot the arm will jump. Some target locations will even have the Baby caught in a cycle.

The rest of the Baby will presumably eventually re-organise to avoid approaching the hot-spot. There are several strategies that will succeed, and the one chosen depends on the learning mechanism.

```
                R*      <- top-level reference
                |
             --------
             | ECSs |   <- multi-level CS (Control System)
             --------
              /    \
             /      \                 R*      <- another top-level reference
            /        \                |
       --------    --------        -------
       | ECSs |    | ECSs |        | ECS |   <- pain reflex ECS
       --------    --------        -------
        /    \      /    \         /     \
      .......................................
       |       \   |       \     Pain    /
       |        \  |        \           /
       +-----------+         \         /      Environment
         |        \           \       /
       Inputs      ------------Outputs
```

So why don't I consider this a cheat too? After all we have hand-coded an ECS to perform a function. Well we haven't had to add a separate learning hierarchy (as per Bill), or had to wire across levels (as in the current Little Baby).

Below are the reasons I think we don't have to add any new features to "force" the Baby to learn the task.

Simplicity:

The pain-reflex is easy to learn by simple means (such as genetic algorithms or random search). We shouldn't need to hand-code such control functions.

Effectiveness:

The pain-reflex is very effective at avoiding the hot-spot. This is accomplished solely by setting the gain high on a simple task.

Stability:

The pain reflex is stable against random re-organisation. Since it is "effective" it very seldom has a non-zero error. (Persistant high local error should probably trigger re-organisation.) Since it is "simple" it has very few weights. This makes it a small target for a random mutation (compared to the rest of the net).

Lastly it is high gain. If there is a random change to an input or output the Baby will thrash wildly. The strong accumulation of local error should quickly cause a benign mutation.

Since the new top-level goal is quite stable the rest of the Little Baby's brain is forced to re-learn.

Now for the proverb. I have realized that one of my original assumptions was wrong. When I first learned PCT I assumed that all the top-level ECSs (ones with fixed references) were also high-level ECSs (far from the environment).

I now suspect that *most* of the top-level goals of an organism are fairly close to the I/O level, and that most of the high-level ECSs are just used to add efficiency to the satisfaction of these low-level goals.

Top-level goals need not be high-level goals.

... Jeff

Date: Tue Aug 18, 1992 8:02 pm PST
Subject: Tumbling Genomes

[from Gary Cziko 920818.2240]

There's a article in the recent _Science_ (vol. 257, 14 August 1992, pp. 884-5, "Possible evolutionary role explored for 'jumping genes'") which I thought at least Bill Powers and Greg Williams would find interesting.

It says:

"If at a certain time the suppressor genes were somehow turned off, the accumulated variation could be unleashed, giving rise to dramatic differences in a population over a short period of time and perhaps even creating changes that could lead to new species."

So here is a possible mechanisms for the genome to "tumble" around a bit when some intrinsic error is high, fitting in nicely with Bill's speculations about evolution as a control process leading to punctuated equilibrium.--Gary

Date: Wed Aug 19, 1992 3:29 am PST
Subject: non-repro blues

(Gary Cziko 920818.1453) --

>to be consistent you should also be leaving messages for yourself on your
>answering machine

We could have a lot of fun with the comparison between CSG-L and and answering machine, but I'll refrain. The analogy is a poor one.

I can identify two pertinent aims:

1. I want a copy of messages I send to CSG-L. I want to be able to go back and reconstruct the context if someone responds in part. This is what is disanalogous to a message on an answering machine. It is analogous instead to keeping a copy of a letter.

2. I want verification that my message was delivered. Experience recommends it. This is analogous to sending registered mail.

I can satisfy (1) by sending myself a cc of each message.

I can satisfy (2) to a degree by some other command (I'll have to look it up so I can turn it off now) that makes the listserver send me a notice of distribution.

However, seeing a report of its action is much less convincing than seeing the results of its action in my mailbox file--much easier to believe that the same results appear in everybody else's mail handler. And the repro command combines 1 and 2 with the setting of one variable.

I'm not trying to persuade anyone else to use repro. Some of us asked how to do this and it seemed we couldn't. Some of us I guess see neither use nor sense in it. But heavens, let's describe it accurately!

Too much time as a tech writer, I guess :-)

Bruce bn@bbn.com

Date: Wed Aug 19, 1992 4:39 am PST
Subject: FYI: refs on machine vision

[From: Bruce Nevin (Wed 920819 08:12:45)]

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

To: machine-learning@BBN.COM, neural-people@BBN.COM
Subject: Some interesting vision references ("Why Progress in Machine Vision Is
So Slow)
From: aboulang@BBN.COM
Date: Sat, 15 Aug 92 11:38:01 EDT

From: defee@tukkasotka.tut.fi (Irek Defee)
Newsgroups: sci.image.processing
Subject: Re: Canny Edge Detector
Date: 8 Aug 92 06:54:58 GMT

In article <1992Aug7.132410.29816@bert.eecs.uic.edu> eddins@bert.eecs.uic.edu (Dr. Steven L. Eddins) writes:
>>The interested hobbyists experimenting with the CAnny edge detector
>>should take a look at the recent papers by Margaret Fleck BEFORE
>>starting to experiment. Just scan IEEE PAMI, I don't have a copy here.
>>Hans
>The references are:
>M. M. Fleck, "Some Defects in Finite-Difference Edge Finders," IEEE
>Transactions on Pattern Analysis and Machine Intelligence, vol. 14, n.
>3, March 1992, pp. 337-345.
>M. M. Fleck, "Multiple Widths Yield Reliable Finite Differences," IEEE
>Transactions on Pattern Analysis and Machine Intelligence, vol. 14, n.
>4, April 1992, pp. 412-429.

Very good reading, indeed. Then, add to the Marr-Hildreth and Canny optimal edge detectors;

J. Shen, S. Castan, "An Optimal Linear Operator for Step Edge Detection," CVGIP: Graphics Models and Image Processing, vol. 54, No. 2, March 1992, pp. 112-133.

But, for God's sake, before developing any new (optimal or not) edge detectors, read please:

R. C. Jain, T. O. Binford, "Ignorance, Myopia and Naivete in Computer Vision Systems," CVGIP: Image Understanding, vol. 53, No. 1, Jan. 1991, pp. 112-117, 1991.

T. Pavlidis, "Why Progress in Machine Vision Is So Slow," Pattern Recognition Letters vol. 13, April 1992, pp. 221-225.

Irek Defee Signal Processing Laboratory,
Tampere University of Technology,
Tampere, Finland, Europe.

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

Date: Wed Aug 19, 1992 7:55 am PST
Subject: Re: PCT & interactionism

[From Bill Powers (920819.0900)]

Penni Sibun (920818, 920819) --

Here is one of our problems:

>when i am driving, snow on the road or rain on the windscreen is
>extremely immediate and concrete.

Here is another:

>i don't have any particular emotional investment in what people
>are made of. i was the one who suggested it's all physics at the
>bottom.

A third:

>you really think that how much work a process takes is determinable by
>introspection? that doesn't sound very scientific to me.

A fourth:

>i was using very gross examples, where i thought the work involved
>would be obvious, that is, it involves muscles, rather than primarily
>neurons.

From my standpoint, these are the same views of behavior that led to behaviorism. It's assumed, I believe, that the Observer can see real "immediate and concrete" reality itself, without interpretation -- just the facts. It's assumed that the world is as it is, and that all you need to know about people is what happens to them and what they do. How they are constructed inside in order that they can behave in that environment as they do is of no interest (and makes no difference). The fact that all motor behavior and all sensory experience is created by neurons is of no importance. In principle, we can understand everything about behavior by watching the interaction of environmental things, events, processes, and situations with organismic activities -- behaviors. Behavior -- what organisms DO -- is to be explained in terms of observable interactions only. It's assumed that the mechanics of behavior will be explained, in the end, by physics and chemistry; there's no need to ask WHAT physics and WHAT chemistry. That is the scientific way of dealing with behavior. If these are really the tenets of interactionism, then interactionism is little different from behaviorism.

Control theory is based on an approach that is basically different from the ground up. It's assumed that perception results from neural activity based on sensory inputs -- that there is no other way to know what is going on outside the organism (and that applies to the scientist as well as to the subjects under study). It's assumed that all Observers must see the world this way, as neural signals standing for a world of which they know nothing directly -- but experienced, of course, as a real concrete external world and a body living in it. It's assumed that all observers act by producing neural signals that activate muscles, and know of their own actions only through sensing of muscle efforts and sensing of the effects on other perceptions. If Observers aren't brains, then at least they get all their experiences and produce all their actions via brains; there's no channel linking awareness to the outside world that bypasses the neural processes of perception. There's no way for them to act other than by sending neural signals to muscles and glands. The entire experienced world, from the most concrete and simple aspects to the most abstract thoughts about them, exists as patterns of neural firing in the brain.

That, of course, is a model. It's a model featuring a device called "the brain," whose internal activities are experienced by a propertyless Observer. It's consistent with our models of physics and chemistry, applied either to the internal parts of the model or to the hypothetical reality outside it. It's consistent with what is known about the physical structure of the body -- biochemistry and neurology. Control theory brings all these models together into a single consistent framework, without claiming any property for a scientist that the subjects don't also have, without claiming that the scientist has any way of acting or knowing the truth that others don't also have. And control theory goes further -- it proposes an internal organization that can account for the way we really observe behavior to work instead of just how it has been imagined to work.

There are, of course, more questions unanswered than answered by HPCT. Many people on this net are trying to answer them, trying out various possibilities, rejecting some and carrying others forward. But behind all these conjectures is a common understanding of the nature of the problem, which is very different from yours. Most of the people who are looking for answers are convinced that just taking appearances for granted and trying to find the rules is futile. They are trying to find a model for the organization of the system that is responsible for both experience and behavior, so that when it is placed in any environment it will behave as real organisms do -- and experience it as at least human organisms do.

There are a few phenomena that PCT has uncovered which are easy to see and which no other theory can yet explain. The main one is that what people do with their muscles is variable, yet the outcomes of the muscle activity are repeatable and resistant to disturbance. From any existing scientific point of view, this phenomenon is counterintuitive and inexplicable.

In the kind of explanatory system you're presenting, this fundamental phenomenon doesn't even appear, because all of your descriptions are cast in terms of the outcomes produced by motor activities -- moving the car here and there on the road, for example. In looking for an explanation of such outcomes, this explanatory system doesn't ask how the behaving system can produce such repeatable outcomes by such variable means. It looks to other aspects of the apparent world for explanations -- those outcomes are just the "easiest" ones that the environment makes possible, elicits, encourages, or whatever. The fact that those outcomes are continuously being disturbed, not aided, by the environment is overlooked: if the outcome is stable, the environment must have made that stability possible.

The fact is that the environment is always working to disrupt that stability. What you describe as an explanation of behavior seems to me more like presenting a series of problems calling for an explanation. The statements you offer as explanations seem to involve more than a bit of magic, and more than a modicum of arm-waving. The arm-waving isn't evident from within the framework you're describing, because there is one question that simply doesn't arise: HOW CAN THAT POSSIBLY WORK? But to anyone who actually tries to make working models of any proposed explanation, the question of HOW is the crux of the matter. If the explanation entails a HOW that is impossible, or that flatly contradicts our other models of reality such as physics and chemistry, that alone is enough grounds to reject the explanation. If the explanation doesn't even consider the question of HOW, then it's not an explanation at all. It's just a description.

You're under a misapprehension about PCT, as Mary pointed out. This model began as a model of human behavior, based on a study not only of real people doing real things, but of physiology and neurology. The clever machine to which the model was first applied was the real thing, a human being. Predictions of human behavior were the first ones made on the basis of the model. The Little Man arm program, the crowd program, the beginning of models of the Beerbug, the interest in Pengi and Sonja, came up only as a way of making the model more detailed, or illustrating the principles of the model in other settings in ways that might communicate to people in other fields.

Whether this communication succeeds or fails depends in large part on whether the recipient of these arguments considers the basic phenomena of control worth some attention, and whether he or she is willing to consider the world as consisting of perceptions rather than an objective world and a perceiver of that world.

Have you tried out the rubber-band demo? I should think that it would offer an excellent comparison of the interactionist kinds of explanation and the control- theoretic kind.

Best, Bill P.

Date: Wed Aug 19, 1992 9:16 am PST
Subject: Drivers, cognition, & PCT and perception

[From Rick Marken (920819.1000)]

penni sibun (920818.2000)

> going from the retinal map to
>seeing objects is not a higher cognitive function--the process of
>perceiving, even in the eye, never mind the brain, is extremely
>complex,

Of course. I assume that the functions that transform sensory inputs into perceptual signals are extremely complex. If the computation carried out by these functions are the "work" of which you speak then we are in violent agreement. This idea of "work" is certainly not new; since at least the mid 60s we've had a pretty good idea that the neural activity that corresponds to perceptual experience is the result of complex functions carried out at many layers of the nervous system. Look at the work of Hubel, Weisel,Ratliff, Letvin, Hartline, -- to say nothing of Ernst Mach. To call these neural computations "work" seems to add little to the notion; the functional computations that are carried out deliver (to me, a least) a continuous, uninterrupted flow of experience. It was in this sense that I rejected the idea that "work" was involved. Sure, there are complex computations, requiring activity in lots of neurons, and certain functions may involve more neural firing than others -- so in that sense there may be differntial amounts of "work" -- but I don't think work is a particularly helpful way to look at perceptual processing. I like to think of it as a functional relationship p = f (s) -- neural firing,p, (that is the perceptual signal) is a function of the sensory signals, s. The function operates continuously over time -- so p is a time varying function of s (where s is a vector, not a scalar). And f() can be quite complex.

> and not at all a matter of passing an uninterpreted signal
>down the line. at least, so say the tons of literature, and our
>resident psychologist, whom i consulted just in case i had missed a
>major shift in the field.

What is an "interpreted" signal? All there are in the brain, as far as we know, are signals. The only possible "interpretation" of the signal is the function, f(), that produced it. The signal is always just a signal; but what it "means" depends on how it was derived. How do you imagine that perception works? Do you think the neuron's carry "interpreted" signals? What are these? Are they like names -- "Here comes a flower perception". Something like that?

>not really. remember, i started out denying that there was such a
>thing as an ``environment.'' and i also denied that there is a locus
>of control anywhere.

Well, then it's going to be real hard to get you interested in control theory. Your assumptions obviate the need to explain the basic phenomenon of control.

>as an aside, i don't think behavior is ``what *needs* to be done.'' i
>think it's what *is* done.

I have spent the last ten years trying to develop demonstrations that would show that precisely that assumption is wrong. The idea the behavior is "what *is* done" is the fundemental assumption of ALL psychologies -- behaviorism, cognitivism, psychoanalysis, psycholinguistics and, apparently, interactionism. You seem to be committed to this point of view -- so I don't imagine that it will be possible to sway you. Most people feel the same way -- at least tacitly. If you are willing to try to see our point of view, I would recommend reading about (and, better yet, doing) what I call my Mind Reading program (described in chap 2 of my book of the same name). The program shows how a person can clearly and unequivocally be "doing" five different things (moving five objects on the screen) but only doing one thing "intentionally". An observer cannot see which of the five behaviors is intentional so any one or all could be called the behavior (doings) of the subject. But, from the subject's perspective, he or she is only doing one thing -- moving one of the objects. The computer, using "the test for the controlled variable", can tell which of these behaviors is being done intentionally -- so it is reading the subject's intention (mind) -- hence, the program does "mind reading".

Have a good time with the interactionism.

Regards   Rick

Date: Wed Aug 19, 1992 10:59 am PST
Subject: Re: interactionism

[Martin Taylor 920819 14:15]
(Penni Sibun 920818.1200) (sorry about the [y] earlier)

> [Martin Taylor 920818 11:50]
> (Penn[i] Sibun 920817)
>
> > there's just tons of literature on
> >perception that suggest that perc. is a *lot* of work; in fact it
> >involves action--perhaps to the degree that p. and a. are inseparable.
> >(``active vision'' is a current buzzphrase.)
>
> This was cast as a criticism of PCT, but in fact it seems to be a reasonable
> description of how PCT says perception works. An ECS that is trying to
> achieve a percept can go only by its error signal.
>
>maybe someone can explain this to me. when you say ``signal'' or
>``variable'' or ``percept'', it has connotations to me of a unified
>thing, like a tone, or a light intensity. but when i look at the
>road, i am not perceiving something like a tone. if you can explain
>how all the stuff my eyes take in can be a single unified thing, maybe
>i won't find it so oversimplified.

(Aside: as a student of language, you won't be surprised to know that I find long postings without sentence capitalization hard to read.)

Maybe the subsequent postings have cleared this up for you, but in case not, here's another attempt.

I don't know how long you have been following CSG-L, but it seems that there is a foundational notion missing. PCT does not deal with the human as a single control system that controls "all the stuff my eyes take in" as "a single unified thing." Instead, there is a vast hierarchy of elementary control systems that make up a hierarchic control system. Each of the elementary control systems (ECSs) do control one unified thing, but that thing (percept) is a function of many inputs. At the very lowest level, the inputs are sensory, but above that they include the results of lower- level perceptual functions in lower-level ECSs. Your perception of being in the middle of the road may be a "unified thing" but it implies the perception of myriads of other variables, such as the intensity of light off that patch of snow to the left, the orientation of the line projected by the centre stripe of the road, the location of the dark corresponding to the windshield pillar, the pressure on your buttocks... You may not be conscious of these other percepts--in PCT the notion of percept is quite distinct from that of consciousness--but without them your perception of being *so far* to the left of the centre line will be less precise than it might otherwise be.

When we say "signal" we usually mean the value of the unified thing that has been derived from all the lower-level complex, or alternatively the reference level with which the percept is compared. "Variable" can be anywhere. I tend to use "CEV" for "Complex Envirnmental Variable" to denote some structure in time and space that corresponds to a percept. One can never be sure that a CEV actually exists "out there" but if our actions control a percept that we feel to be based on a particular CEV, then we have some evidence that the CEV is a useful perceptual construct. It is "real" for us (though perhaps not for anyone else). Internally, percepts, reference signals, errors, output, etc. are all "variables," and there are probably many more. "Democracy" could be a perceptual function that has a particular level in specific curcumstances (very low level these four days in the USA, I'm afraid, for me), but if there is a CEV corresponding to my (unitary) percept "democracy," I'm pretty sure it won't be exactly the same CEV that affects your perception of "democracy."

Another word to be careful of in PCT terms is "action." Actions are observable but hard to interpret. An observer sees hands and arms moving, but if those actions were duplicated on another occasion, it is highly unlikely they would have any useful effect. Actions affect CEVs, and the effects lead to changes in the percepts associated with the affected CEVs. But unless the observer knows what the CEV *is*, the actions are meaningless. And to know what the CEV is, is to know the intentions of the actor.

You are talking, for example, with a few other people, and one gets up and opens the door. You see the arm and hand manipulating the doorknob. But the real question you can ask him or her is "Are you leaving, or are you leaving?" Your two alternatives are opposites, couched in exactly the same words. You want to know whether he or she is finished with the discussion or is coming back after a short but necessary break. You need to know what CEV is being acted upon. Is it one related to the interest value of the discussion, or is it related to internal physiological states? The actions are the same, but the implications are very different. And so it is with your question. Two almost identical actions (word strings with intonation) have meanings associated with the CEVs you see as possibly related to controlled percepts in the person opening the door. All unitary, all complex in construction.

Anyway, the root misunderstanding expressed in your quoted paragraph seems to be based on the idea that there is a single unitary percept that we control, whereas in PCT terms, each ECS controls a single unitary percept, but there are myriads of them involved in even very simple real-world activities such as driving.

Hope this helps Martin

Date: Wed Aug 19, 1992 11:11 am PST
Subject: Re: Drivers, cognition, & PCT and perception

[Martin Taylor 920819 15:00]
(Penni Sibun 920818.2000)

>i am deeply suspicious of hierarchical organization from my experience in
>modelling language (as well as other behavior). language is supposed to have
>all sorts of hierarchy in it (eg, a text is composed of paragraphs are
composed >of sentences are composed of words) and i just don't think this is a
very useful >way of looking at language, and am working to show this.

I would be very interested to know more. I have evolved in exactly the opposite direction. My theory of Layered Protocols seems to me to hold promise of clarifying many of the difficult problems of language, and I think that arguments from information theory and feedback stability theory almost demand that dialogue be conducted in a hierarchic way. If you have counterarguments, I'd be most interested in them.

But if all you mean is that one cannot describe language as the simple combination of unitary elements of different sizes, then I totally agree. Your use of the word "composed" suggests that this may be your meaning. All the elements, of whatever size, change their functions dramatically depending on situational and linguistic context, and the contextual effects can be very long-lasting. Language is not a concatenative process.

I do hope you can learn enough about PCT to appreciate what it is doing. It took me over a year of working with, and arguing in, CSG-L to realize its beauty and power, and even now I doubt I have seen more than a glimmer. I have come to feel that it really is the Newtonian revolution that physics had 300 years ago and that psychology needs. After 35 years as an experimental psychologist and sometime psycholinguist looking for some coherence in the field, it's quite a revelation to find that one really can do more than produce little models that work for specialized micro-fields of enquiry.

Martin

Date: Wed Aug 19, 1992 11:36 am PST
Subject: Locus of Control, Behavior

[FROM: Dennis Delprato (920819)]

IN SUBJECT: Drivers, cognition, & PCT and perception

>>(Rick Marken (920819.1000))

>penni sibun (920818.2000)

>not really. remember, i started out denying that there was such a thing as an ``environment.'' and i also denied that there is a locus of control anywhere.

>>Well, then it's going to be real hard to get you interested in control
>>theory. Your assumptions obviate the need to explain the basic phenomenon
>>of control.

On the conundrum of locus of control: It is my understanding that HPCT takes the traditional question of locus of control (it's inside the organism--mentalists/cognitionists; it's outside the organism--most behaviorisms; it's both inside and outside--cogni- tive behaviorists) as poorly put. Rather one part of the control does not control any other part. The entire system IS a control system. Another conundrum, behavior, enters in here (below). Furthermore, control is not an "it" that one can locate anywhere.

>as an aside, i don't think behavior is ``what *needs* to be done.'' i
>think it's what *is* done.

>>I have spent the last ten years trying to develop demonstrations that
>>would show that precisely that assumption is wrong. The idea the behavior
>>is "what *is* done" is the fundemental assumption of ALL psychologies --
>>behaviorism, cognitivism, psychoanalysis, psycholinguistics and, apparently,

>>interactionism. You seem to be committed to this point of view -- so I don't
>>imagine that it will be possible to sway you. Most people feel the same
>>way -- at least tacitly. If you are willing to try to see our point of view,
>>I would recommend reading about (and, better yet, doing) what I call my
>>Mind Reading program (described in chap 2 of my book of the same name).
>>The program shows how a person can clearly and unequivocally be "doing"
>>five different things (moving five objects on the screen)

The five different "things" are five different behaviors, but none are of interest to the (HPCT) psychologist. They are not of interest because they are structural, topographic, or physical behaviors. To scrutinize them is to learn what the person is doing--physically, not psychologically. I am getting at what I thought was understood in HPCT, but perhaps needs reemphasized. This is that there is behavior and there is behavior. In the most elementary sense, it is crucial to distinguish between behavior-1, or physical behavior (the behavior of the moon, of molecules, of stones), and behavior-2, or psychological behavior. Most psychology up to the present (as exhibited in behaviorism, information-processing theory) has usually taken behavior as behavior-1. Many thinkers have come close to making the break--drawing a firm line between behavior-1 and behavior-2, but have not been successful. Freud, Brentano, James, and Skinner, among others, have almost made convincing cases, but it is clear that we continue confusing behavior-1 and behavior-2.

>>but only
>>doing one thing "intentionally".

THIS is recognizing behavior-2. Intention, purpose are not found in the behavior of the moon, in the behavior of atoms.

>>An observer cannot see which of the
>>five behaviors is intentional so any one or all could be called the
>>behavior (doings) of the subject. But, from the subject's perspective, he or
>>she is only doing one thing -- moving one of the objects.

Translation: In the absence of certain strategies (perhaps "the test"), an observer cannot except speculatively identify what the psychological event here is. HPCT provides a theoretical framework and methodology for understanding psychological events.

One additional point on behavior. To state that behavior is the control of perception is not equivalent to stating that behavior (an "it"?) controls perception. I think that sometimes the latter is implied in certain statements beginners to HPCT might make to themselves or to others.

I hardly speak for HPCT. Consider me a student.

Dennis Delprato
Department of Psychology
Eastern Michigan University
Ypsilanti, MI 48197 U.S.A.
Psy_Delprato@emunix.emich.edu

Date: Wed Aug 19, 1992 12:20 pm PST
Subject: Re: PCT & interactionism

[From Mary Powers 9208.19]
(Penni Sibun 9208.18)

>the road ... is constantly changing ...

yes. But it is your perception of the changes that matter to you. A foot of snow on it is a big disturbance, and you alter your driving actions accordingly. A freshly painted yellow stripe may be noticed, but it's not a disturbance.

>it's an artifact--somebody designed and built it--and these are important things >to know about it.

Important to who? It sounds like you are controlling for - have a high reference signal for - being aware of these things. I'll bet 99% of the people who drive down the road don't give a damn. This tells me something about you, not about any intrinsic importance.

>I am deeply suspicious of hierarchical organization.

So am I. Word/sentence/paragraph is as pseudo-hierarchical as cell/organ/organism/group. They are all at the category level.

>I really can't imagine how a custom can be a program ...

I misspoke somewhat. I mentioned rules of the road as principles. Stop at red lights is a sub-principle. This sets the reference signal for the following program: If red, then stop. This sets the reference signal for the sequence: take right foot off accelerator, put on brake and push. Each part of this sequence is accomplished by reference signals to lower and yet lower levels, the last one being the actual effort of the muscles. All this brought on by noticing a brightness, which is processed as red, which is observed as being part of a traffic signal, which is in my lane, which I am approaching ... which I compare with what I want to be doing when I see a red light, conclude that I am not doing it, and proceed to initiate the program above.

>how does an individual introspect about it?

I think you introspect by running it through imagination (see Behavior: the Control of Perception - by the way, if you're trying to get all this through the net without reading the basic texts you're asking people to repeat an awful lot of work that's already been done)

>how are customs passed on and maintained?

Ask a PCT anthropologist. I wish there was one. Children don't come equipped with the full hierarchy. As it develops they try setting reference signals similar to the ones around them. They get love and approval for this, which satisfies inborn reference signals for being loved and approved of.

>how are they recorded?

In memory, whatever that is. Ask a neuro-chemico-physio-psychologist. He probably doesn't know either.

>how are they incorporated into things like building (better) roads?

Somebody in the highway department gets a lot of complaints about a dangerous intersection. He is controlling for minimizing traffic accidents. From his experience he knows that adding a traffic light and some left turn lanes will do this, so he orders it done (budget permitting).

It's important to realize that his order is not setting a reference level for the people doing the job. For them his order is a perception, and is processed through the nervous system of a subordinate like any other perception, where it is compared with reference perceptions already there. It is likely that the subordinate will have a reference signal for the order that matches it - this is what I get paid to do - and an error - it isn't done yet - and do what has to be done. There's a half-baked idea floating around that organizations are super-organisms, with the guy at the top setting reference signals and the people below him being inputters and outputters. Not PCT, for sure.

Mary P.

Date: Wed Aug 19, 1992 1:30 pm PST
Subject: Re: interactionism

[From Rick Marken (920819.1400)]

Martin Taylor (920819 14:15) to penni sibun

>PCT does not deal with the human as a single control system that controls "all >the stuff my eyes take in" as "a single unified thing."

Nice post Martin. I think you honed in on a common misconception about PCT. If I didn't already understand your point I know that your nice discussion would have helped me towards that understanding.

Dennis Delprato (920819) says:

>On the conundrum of locus of control: It is my understanding that HPCT takes >the traditional question of locus of control ...as poorly put. Rather one part >of the control does not control any other part. The entire system IS a control >system.

Well, I'd be inclined to say that control theory does place the locus of control squarely inside the control system (not in the environment). Variables in the environment are controlled (maintained in reference states against disturbance) and this occurs because of the nature of the control system itself (closed loop, negative feedback and, especially, HIGH GAIN); not because of the nature of the environment. So whoever said that the locus of control was inside the organism (like the cognitivists, maybe) was right; PCT just shows why this is the case and what it imples about the nature of behavior (it is controlled perceptual experience; I like to say it means that behavior is a subjective, not an objective phenomenon -- the cognitivists never got THAT part).

> there is behavior and there is behavior. In

>the most elementary sense, it is crucial to distinguish between behavior-1,
>or physical behavior (the behavior of the moon, of molecules, of stones),
>and behavior-2, or psychological behavior. Most psychology up to the
>present (as exhibited in behaviorism, information-processing theory)
>has usually taken behavior as behavior-1.

Nice distinction. Yes, PCT is interested in behavior-2, and sees behavior-1 as just a side-effect of behavior-2, interesting to an observer, perhaps, but usually irrelevant to the behaving system.

>In the absence of certain strategies (perhaps "the test"),
>an observer cannot except speculatively identify what the psychological
>event here is. [In the mind reading program]

Correct. And the main reason is because the disturbances (which are continuously influencing all 5 "behavior-1" behaviors are invisible. So it is truly impossible to tell (just by looking at the behavior of all 5 objects -- or, I might add, by doing any kind of analysis of their behavior-1 behavior) which is being moved intentionally. In fact, the idea that ANY one is moving intentinoally would not even occur to most observers; all the objects just move around in an arbitrary pattern. The fact that one of the visible behavior-1's is the result, in part, of a person's active efforts to resist a disturbance is completely invisible -- because you can't see the disturbances. The situation is similar to the one we are in when we just watch people behave; all that behavior looks like behavior-1 because the disturbances are invisible so the resistance to disturbance is invisible. Control theory suggests that some of the events we see as behavior-1 may be the result of behavior-2 -- ie. control. We can tell which behavior-1's are the result of behavior-2 by disturbing a behavior-1 variable and looking for lack of (or considerably reduced) effect. The mind reading program does this because it can "see" the disturbances that the observer of the behavior of the objects cannot see; so the program can determine which disturbance is being systematically opposed.

Best regards Rick

Date: Wed Aug 19, 1992 1:52 pm PST
Subject: Misc catching up

[From Bill Powers (920819.1500)]

Bruce Nevin, Gary Csiko, recent posts.

You're right, I'm wrong. With 100 people on the net, asking for a copy of your own posts increases the activity 1 percent. Feel free.

My mainframe at Ft. Lewis, by the way, wouold put a copy of each transmission in my directory if I told it to. As I edit and read mail offline, I have copies of everything anyway.

Catching up: Gary Cziko (couple of days ago)

"Computing output" has suffered the overgeneralization that all slogans seem to create. It's short for "computing the exact output that will create a predetermined outcome when it's executed."

This is different from computing an output in real time that will tend to make the present-time error get a little smaller.

Also, "computing" an output is a modeler's way of looking at things. A modeler would say that a transistor amplifier computes (outp[ut signal) = k*(input signal). Of course it does no such thing; it isn't computing as a digital computer would compute. It isn't even computing as an analog computer would do. It's just behaving according to physical principles, which we REPRESENT as a computation in a model. This leads to a slight digression.

The same is true in most of the nervous system. Most of the levels of control, below the program level, don't literally compute anything. Perceptual signals at one level depend on sets of perceptual signals at lower levels because of the way the signals are combined in neural nets. Only at the level where we explictly apply analog or discrete methods of computation is there any literal computing going on.

This is one reason that talking about lower levels as if they worked in terms of language is a mistake. A relationship-controlling system doesn't say "The car is too far left; I'd better turn the wheel to the right." The perceptual signal is compared with the reference signal by inhibiting part of it; the leftover part, the error signal, passes to lower levels via output amplifiers and weighters, adjusting the lower reference signals and eventually the amount and direction of effort -- all without any commentary or reasoning.

Of course at the same time, the driver's higher-level systems may be conducting imaginary conversations about what's being experienced at the lower levels. "Wow, I just missed that baby carriage. I'd better give it some more gas going around the next curve to improve the traction. Good thing this car is so heavy; it stops faster than a light car. Good thing I have fast reflexes." Blah, blah, blah, all sorts of opinions and misinformation that have absolutely nothing to do with how the car is being controlled or what skills are actually being used. The verbal obbligato is not part of the control processes going on, unless it's specifically about resetting reference signals: "Oops, I forgot to pick up my wife, I have to go back."

By talking, we can make behavior seem a lot more complicated than it is.

Dag Forssell (920817) --

>I will grant you that everything is variable, but we have been using the term >variable to signify that feature of the environment that we want to control.

What you mean we, white man? I've experimented with lots of formal ways of using terms. For while, a variable in the environment was a quantity, not a variable, while inside the nervous system a variable was a signal. Now I just use variable in its mathematical sense, something that varies or can vary. The largest category is variable, and we divide it up into smaller classes as needed: variable, input. Variable, disturbing or independent. Variable, controlled. Variable, output. Variable, reference. Variable, error.

I suppose that for introductory teaching it's a good idea to stick to one terminology until the ideas are firmly grasped. But after that, I think it's better to vary the language so people will always have to think about what it means, and not turn the words into slogans or just memorize them. You should be able to explain control theory many times without ever repeating your words. When you can do that, you know you understand what it means.

John van Loon (920817) --

>This term I will be building an autonomous vehicle. I will try to
>organize it to accept PCT concepts and hopefully it will serve as a
>viable testing ground ...

Good! You'll learn more about control systems by building one than from a hundred posts.

I suggest thinking in terms of perceptions first. Give your car some photocells, maybe some thermistors to detect heat, whatever else is easy to do and cheap. If you're using a digital computer, get an A-to-D board so you can handle analog inputs. Then set reference signals, compute errors, and use the errors to adjust direction and velocity.

One caution. There's a temptation to connect the error signal to the steering wheel to set the angle of the front wheels. In steering, the direction of motion changes as the integral of the wheel angle. You may have to add some first derivative to the error signal (error + d/dt(error) to stabilize the steering.

A tip: remember to design your system as if you're inside it looking through its sensors, knowing only the signals from the sensors.

Good luck.

Avery Andrews (920819.1231) --

>PCT in fact depends very much on the fact that many (in fact most, by an
>overwhelmingly large margin) features of the environment can be leaned on.
A >steering subsystem would for example depend on all sorts of facts about
roads >and carparts about which it knows nothing, as well as upon the workings
of the >lower-level control-systems it works through.

I don't know what this image of "leaning on the environment" is supposed to mean. The environment is what it is, it changes when you act on it as it changes. It often changes all by itself, causing you great trouble. Your goals are usually stated in terms of perceptions of the environment that you've found pleasing or beneficial to control. But just having a goal doesn't make the environment be that way. You have to do something to the environment, and when you do, you find out you have to do things a little differently each time you try to control for the same thing. Often a lot differently. The environment isn't very cooperative; it isn't very reliable. Right now, Mary's priming the back deck to keep it from rotting out. I'm still trying to get a piece of video software to install right.

Most of the things in the environment that seem repeatable and reliable are kept that way by someone's close attention and hard labor. Left to itself, the environment would soon be overgrown with weeds and full of holes. Like my yard, which I am allowing to have its own way, more or less. To get a view of what nature is really like, take a walk through some undeveloped and unmaintained woods.

It's a good thing there are some regularities in nature. If there weren't we couldn't control at all. But there are far fewer regularities than seem to be present.

RE: beerbugs

>What goes on here is that if the bug is hungry enough (the energy level is far >off enough from its reference level) and over food, it will start eating, but >being eating also excites the feeding controller, so it keeps eating even when >its energy level is no longer low enough to successfully trigger eating.

This arrangement comes from thinking of eating as being "triggered" like a flip-flop. If eating becomes possible, a control-system type model would simply eat until the reference signal was matched, and then it would stop. There can be two levels to this: an ingestion system with its own reference level for "fullness", and a longer-term nutrition system that monitors a more slowing-varying variable that is gradually changed by eating and not eating. With the right adjustments of parameters, you get bouts of eating separated by not eating. The system is unstable, the more so if it's very hungry (nutritionally). I must resurrect that operant-conditioning model before it gets lost. The "scalloping" you get from using too large a reward for each reinforcement is a simple oscillatory phenomenon, which the model shows without any special provision for it.

The motto is, don't look for fancy explanations and models when a simple one will do just as well.

Best to all, Bill P.

Date: Wed Aug 19, 1992 3:55 pm PST
Subject: Re: Misc catching up

[Avery Andrews 920820.0945] (Bill Powers 920819.1500)

Leaning:

I agree with almost everything you say about the environment (and so do Chapman and Agre, who say that the natural world is much more dangerous than the artificial one we make for ourselves with culture). The target of the `leaning on the environment' slogan are the `planning weenies', who really seem(ed?) to think that a competent Agent ought to calculate the results of its actions from first principles. As to `how many' regularities there are in nature, I doubt that this is a question that can be profitably argued about. At a very minimum, PCT depends on there being consistent correlations between the direction of gross effort and the direction of change of net result, & a continuous function relating the magnitudes of these two quantities. One is thus `leaning' on everything that causes these correlations to exist, without having to know anything about these underlying causes or do any computations involving them (like a conventional Planning Agent would have to, if what I read about them in Beer, Chapman, and Agre is true (I haven't really gotten around to Brooks yet)).

Beerbugs:

I'm still confused. I have on the page before me a diagram of a circuit that ought to work, and is supposedly based on the actual wiring of a real critter, and which I believe I described accurately. Even if Aplysia doesn't actually work that way, I don't see any reason why some other critter couldn't, & I don't (yet) see how to apply your description to the actual circuit.

But maybe the light will dawn....

Avery.Andrews@anu.edu.au

Date: Wed Aug 19, 1992 4:26 pm PST
Subject: Re: Misc catching up

[Martin Taylor 920819 2010]
(Avery Andrews 920820.0945) A quick response, indeed!

I imagine Rick will respond to this, but in case not:

> At a very minimum, PCT depends on there being consistent correlations between >the direction of gross effort and the direction of change of net result, & a >continuous function relating the magnitudes of these two quantities.

Only at a very high level, if there. A hierarchic control structure can accommodate a world the reverses the direction of effect of actions. A reorganizing system can learn to do it more efficiently if the reversal is maintained, but a fixed system can be developed that can invert the sign of its output if the error is increasing. Bill posted a note a little while ago showing how a 2-way ECS made from a pair of one-way ECSs could be provided with a signal that affected its gain. The same principle could be used in building a reverser control, I think. Anyway, humans have no problem with such reversals, as Rick (with Bill? I forget) showed.

I think that consistency is probably required over a time commensurate with the inverse bandwidth of the feedback loop concerned, but that's only hazarding a guess. There is no requirement for a continuous function, either. But I think there probably is a requirement for some measure of closeness between percept and reference.

Analysis is simpler when the relations are continuous and consistent, and we usually talk about that kind of system.

Date: Wed Aug 19, 1992 5:39 pm PST
Subject: Re: Misc catching up

[Avery Andrews 920820.1115]

(Martin Taylor 920819 2010)

> A quick response, indeed!

Yes - I've even found myself with plane reservations leaving Hawaii the day before I got there.

>Only at a very high level,...

Corrections accepted, but what exactly is `inverse bandwidth'? If I had been being more careful I would have said something to the effect that the correlations between direction of effort and direction of change in result must persist for `reasonable' time periods, that is, not changing too often. But I don't have the conceptual tools to know how to be precise about how often is too often.

Avery.Andrews@anu.edu.au

Date: Wed Aug 19, 1992 7:34 pm PST
Subject: bit.listserv.csg-l

[from Gary Cziko 920819.2155]

Martin Taylor 920818 18:30 said:

>(I don't know the usenet name for the extension of CSG-L).

I apologize for being delinquent in my planned monthly informational posts to CSGnet which would include the Usenet (NetNews) name for CSGnet.

Our newsgroup name is bit.listserv.csg-l. This may not be available on all systems which include the sci. groups, but in theory can be added to any Usenet node. For CSGnetters with access to Usenet (NetNews), this may be a more convenient means of interacting with CSGnet (it is more convenient from my perspective since I don't get returned mail from bad addresses).--Gary

Date: Wed Aug 19, 1992 7:44 pm PST
Subject: Re: PCT & interactionism

(penni sibun 920819)

i think we've all gotten rather tired of this thread so i will summarize and i think drop it, unless someone has a specific question. another reason to drop it, as i've mentioned, is that i have repetitive motion injury troubles, and really oughtn't to be typing such screeds. (and since the problem is worst in my right pinkie, capitals are a luxury i can't afford, unfortunately--sorry!)

we started out by someone suggesting that interactionism is behaviorism. i objected, and said that pct sounded cognitivist to me. i think these two views are telling: they certainly tend to pit us against each other, because of the historical adversary bet. beh. and cog..

i'm sorry if it sounded as though i thought pct started out building creatures and generalizing to people: i meant that to *me* pct looks quite plausible for building creatures and rather less plausible for explaining non-built organisms.

i think the importance of the brain is pretty clear, but i don't think the brain is the be-all and end-all of organismhood. what about plants and bacteria and viruses and mitochondria? what about anthills and lichen? in creatures with a brain, what about the spinal chord and sensory organs and all the afferent and efferent neurons? as far as the brain itself goes, no one really knows how it works. what about hormones? what about recent research suggesting that the brain may communicate w/in itself via hormones?

i think it's inportant to try to model the brain; i think it's also important to be open to the idea that that might turn out not to be the best thing to model.

i don't think i can succintly explain active perception; i'm sorry about that. as i say, there's a literature; it's an active area of cognitive science these days.

> [From Bill Powers (920819.0900)]
>
> Control theory brings all these
> models together into a single consistent framework, without claiming any property for a scientist that the subjects don't also have, without claiming that the scientist has any way of acting or knowing the truth that others don't also have.

i find this opposition odd (you've brought it up before), because i don't know anything i or agre or chapman or preston or whoever has said that has claimed there is a privileged Observer. rather the opposite: it seems that a point of view that suggests that organisms aren't privileged cannot coherently claim that a particular organism is privileged.

> In the kind of explanatory system you're presenting, this fundamental phenomenon doesn't even appear, because all of your descriptions are cast in terms of the outcomes produced by motor activities -- moving the car here and there on the road, for example.

i *thought* all my descriptions were in terms of things like snow and road reflectors and what i've overheard people say. in my descriptions, motor activities played a very minor role. in fact, i don't believe i once mentioned moving anything anywhere.

> Have you tried out the rubber-band demo? I should think that it would offer an excellent comparison of the interactionist kinds of explanation and the control-theoretic kind.

no; in what medium is it?

thanks, y'all, for the discussion.

cheers. --penni

Date: Thu Aug 20, 1992 6:58 am PST
Subject: Re: PCT & interactionism

[Martin Taylor 920820 10:30]
(Penni Sibun 920819)

Sorry about your repetitive motion injury. I do hope you are getting medical attention for it and not just trying to cure it by yourself (end parental mode).

>i think the importance of the brain is pretty clear, but i don't think the brain >is the be-all and end-all of organismhood. what about plants and bacteria and >viruses and mitochondria? what about anthills and lichen? in creatures with >a brain, what about the spinal chord and sensory organs and all the afferent and >efferent neurons? as far as the brain itself goes, no one really knows how it >works. what about hormones? what about recent research suggesting that the >brain may communicate w/in itself via hormones?

I can't speak for anyone else on CSG, but I have never thought of identifying the control hierarchy with what goes on in the brain alone. Bill has made it clear that other kinds of organisms can be control systems, although he sometimes seems to obscure the claim by saying that the hierarchy is implemented in the CNS. We have recently discussed the control system of Paramecium, a single-celled organism without a nervous system, which can move in a temperature gradient to control its temperature under hypoxic conditions. Any factor that can have an effect that can be sensed by the generator of the factor can be a component of a control system. One presumes that early almost-living things had control systems that were entirely chemical, and we know that cells in present living systems have many such control systems.

In CSG-L we usually talk about human sensory-motor control systems, and we presume (as much for convenience of discussion as for any other reason) that the communication among ECSs happens through the neural channels. But we have also discussed the possibility that a single neuron could implement several cooperating ECSs, based on the different functions of synaptic junctions near to and far from dendrite roots.

Trees affect the chemistry of the ground near them, altering their growth possibilities. They are control systems if the loop has negative feedback, so as to maintain an optimum local ecology for the treee. If it doesn't, they are not. There's no bias in the theory toward or away from any kind of life. We just are naturally more interested in human life and "intelligent" function than we are in the life of trees. Other scientists may have a different bias. But where control exists, it is important in stabilizing something important to the life of the organism, and should be considered in studying that organism. We deal in psychology, mainly. It's only a part of what PCT offers.

Martin

Date: Thu Aug 20, 1992 8:03 am PST
Subject: Re: PCT & interactionism, Blindmen II

[From Rick Marken (920820.0900)]

Martin Taylor (920820 10:30) --

Gee Martin, how can I add to your comments when you say it all so well.

>I can't speak for anyone else on CSG, but I have never thought of identifying
>the control hierarchy with what goes on in the brain alone.

Absolutely! I use "brain" to mean nervous system (including hormonal systems and whatever else) -- and my main interest is in understanding human beings (a group from which I now exclude members of the republican party) who are typically endowed with brains and nervous systems. But I am also interested in control in systems that have no brain or nervous system -- E. coli, for example (and members of the republican party). E. Coli, by the way, is another good example of a system that controls although there is no correlation between actions (tumbles) and the result of those actions (direction of movement after tumble) -- but there is still control.

Note-- I have re-written (slightly, but brilliantly) the "Blind men" paper as a commentary on an article by Herb Simon that appeared in a recent issue of Psychological Science. It is relevant to our current discussion of "interactionism". Simon said that we have all kinds of great computer models of behavior -- so we understand behavior better than we think. My commentary said "that depends on what you mean by "behavior", Nobel man". I really only added a couple paragraphs to the old "Blind man" paper. But if anyone is interested, I will be happy to post the new version. I've already submitted it to Psych Science but I'd still appreciate comments.

Best regards Rick

Date: Thu Aug 20, 1992 9:11 am PST
Subject: Little Baby & Reorganization

[From Bill Powers (920818.2300)]

Jeff (910818) --

You are running into the same conceptual problems that eventually led me to propose the reorganizing system as I did.

In this diagram:

```
                R*      <- top-level reference
                   |
                --------
                | ECSs |   <- top-level ECS
                --------
                /     \
               |       --------
      direct ->|       | ECSs |   <-untrained ECSs
      input    |       --------
      to top   |       /     \
               |      /       \
      ...................................
                \   /          \
                Input          Output     <- environment
```

... you have both levels of system controlling the same input variable. If you simply connected the output of the higher system to the output below the line, you'd have a control system.

In effect, you're trying to introduce a teacher that already knows how to control the input, and it's teaching another system to do it the same way. This is the problem that all modelers of self-organizing systems eventually have to face. The problem is that reorganization is goal-directed: it must aim to build an organization that is good for the organism. Its goals, therefore, must already take into account the "value" (to the organism) of the variables it senses, and its actions must build behaving systems that control the environment in a way that satisfies these values. But the untrained ECS is not yet an ECS, so it doesn't know what to sense, or how any particular output will affect what it senses. It doesn't know what a "good" target setting of the sensed variable would be. So it seems that the reorganizing system already has to be able to control the same sensed variable that the untrained system is going to be controlling, and it must know the reference level that is best, and it must know what outputs are needed to effect control. If that's true, why doesn't the higher system simply actuate the outputs to control the variable? If it already knows what to do, why doesn't it just do it?

In the second version, you have a control system that avoids touching a hot place. This system is built in, and its reference signal is fixed. You say The rest of the Baby will presumably eventually re-organise to avoid approaching the hot-spot. There are several strategies that will succeed, and the one chosen depends on the learning mechanism.

The fixed control system is an isolated parasite in the hierarchy: it opposes any attempt by the rest of the system to use the same position control system to touch the hot spot. It does so by creating a conflict, which arises from the fact that it acts by using the same muscles that all other control systems must use. Its reference signal, however, does not come from the hierarchy; it is fixed. This system is therefore like a separate organism living inside the hierarchy. Its actions, being unrelated to activities in the hierarchy, constitute a disturbance, how powerful depending on its loop gain. The rest of the hierarchy will not learn to avoid the hot spot; it will (if it can) learn to overcome the disturbance. You have given the rest of the hierarchy no reason to avoid touching this spot; that is the business of the built-in circuit. So there is no basis for the rest of the hierarchy to avoid trying to touch the hot spot. You have put the recognition of the value of touching the hot spot into one low-level system; to the rest of the hierarchy, pain is then just another sensory signal, and there's no _a priori_ reason to hold it at any particular reference level.

You still haven't explained reorganization: the parasitic system doesn't reorganize. You have only created a situation in which reorganization in the hierarchy would be necessary if the parasitic system interferes with control by the rest of the hierarchy. So we are still left with the question of how the rest of the hierarchy reorganizes, and what the basis for and mechanism of such reorganization would be.

The problem with any built-in system or teacher is that you end up having to put all the intelligence needed for control into the built-in system. For instance, if Baby's finger reaches out and touches a hot spot, the "high-level" system has to understand that the remedy is to pull the arm back. If the arm pulls too far back and the elbow touches another hot spot, the same system has to understand that now the remedy is to move the arm forward. If the arm is immobilized by an obstacle, the system has to know that the other hand must be used to knock the hot object away. So this system has to know about relationships between painful sensations on all parts of the body and the effects of tensing various muscles attached to the skeleton. Before you're through designing this high-level system, you will have built in the entire control hierarchy -- which then, presumably, copies itself into the untrained part of the CNS. But in that case, why bother with the untrained part? You're describing a system whose behavior is entirely instinctive from birth.

There may be some built-in control systems low in the human hierarchy, at birth. There probably are. But few of the "instinctive" behaviors we are born with last long; most of them are reorganized away very early in life, or turn into low-level control systems whose reference signals are set from above. It's been known for decades that the ENTIRE nervous system is highly plastic, even at the level of the spinal cord, the so-called "reflexes." We are not a mass of built-in reflexes with higher systems trying to behave around them. As we learn the higher levels of control, we lose the built-in behaviors, or their reference signals become variable.

Low-level systems with fixed references are disconnected from the hierarchy. The more of them there are, the less freedom there is for the higher systems to behave. Remember that ALL systems have to act by using the same set of about 600-800 spinal control systems. Those control systems must be accessible, via their reference signals, to ALL systems that need to use them to generate actions. And that means all systems that produce behavior.

Another problem with considering things like pain a "high-level" controlled variable is that you then have to consider pain to be similar to things like morals or patriotism, which can lead a person to suffer pain deliberately, like G. Gordon Liddy. But if you delegate the control of pain to a hierarchically low level of control, it becomes impossible to explain why pain has a negative value at higher levels or how the higher levels can deliberately accept pain.

Neither of the approaches you outline above can explain real behavior. If there's a built-in control system for avoiding hot spots, how hot does the spot have to be in order to prevent touching it? If approach to the hot spot, and withdrawal from it, are under control of an automatic system, how is it that a person can hold the hands at just the right distance from a fire to keep them nicely warm? How is it that a person can take hold of a hot object and quickly move it to a safer place, suffering pain for an instant in order to avoid some higher-level problem like melting the handle of a pot?

Even more important, how is it that a person can learn not to touch things that are ARBITRARILY hot? The burner of an electric stove that has been turned off a few minutes ago. A light bulb that looks perfectly ordinary but was just turned off. A soldering iron that doesn't glow. A dark rock or a sandy beach out in the sunshine. There's no way that a built-in system can know what aspects of the environment are likely to give rise to pain. It doesn't have enough complexity in its sensors, and even if it does, those sensors couldn't have evolved to understand an environment that is infinitely variable, and in which relationships exist which have never existed before in evolutionary time. The real test that a reorganizing principle must pass is that it must produce an organization that is good for the organism in an environment where there is no _a priori_ reason for any particular sensed state of affairs being either good or bad for the organism. The built-in value of things like pain must guide reorganization at ALL levels, including the highest.

Reorganization has to be based on the control of variables that have _a priori_ value to the organism. Nothing in the external world, as sensed, can have any such value, not for a human being. The external world is simply too variable. So reorganization can't be based on the state of the external world. It must be based on the state of the internal world.

The reaction to pain must be based not just on the pain signals but on the REASON FOR WHICH pain signals are not good for the organism. That is, the value of pain does not lie in the sensory signal; it lies in the state of the organism that gives rise to the sensory signal. It is possible, in short, that there are _a priori_ aspects of the state of the organism that serve as indicators of viability. These are the things that a reorganizing system must sense. In effect, the reorganizing system must sense and control for _viability itself_ as indicated by certain critical variables that it monitors.

In keeping with the same principles that apply in the hierarchy, the output effects of the reorganizing system do not have to employ outputs that are the same as its inputs. We control sensed force not by acting on the tendon receptors, but by sending signals to the muscles. We control the appearance of objects in the world not by somehow creating a direct effect on the objects, but by setting reference signals for control systems that vary the angles at the joints in arms, wrist, and fingers. We cause one controlled variable to change as a means of controlling a different one that is dependent on, but physically different from, the one we use as output.

This is my principle of reorganization. If the pH of the blood stream is too high, the reorganizing system does not react to the drop in pH by increasing pH directly; it doesn't know how. It does so by altering systems that control for sensory variables. Eventually it will alter control of just those sensory variables that indicate states of the environment that bear on the state of blood pH. The organism as a whole learns not to hold its breath too long, or to avoid doing things, like diving too deep and too long under water, that have an adverse effect on blood pH. The reorganizing system itself knows no more about HOW various behaviors affect blood pH than it understands the physics of the environment that make one object's behavior depend on another object's behavior. It doesn't care about those external processes. All it cares about is blood pH. It will accept any behavioral organization, however bizarre, that proves to bring the pH back to its proper reference level. By "accept" I mean simply that it will stop reorganizing the behaving systems.

An example. If you were building a photoelectrically recharged battery-powered robot to roam around free indefinitely, one of the internal aspects its reorganizing system would have to sense would be the state of the battery voltage. Another might be internal temperature of its circuit boards. It would not have an intrinsic reference level for being in sunlight or even for avoiding places where sunlight is unavailable. To know about such things, the reorganizing system would have to have exteroceptors and complex enough perceptions to be able to know where it is in the environment, and what sorts of places would tend to be shadowed. You'd end up putting most of the behavioral hierarchy and its intelligence into the reorganizing system.

By sensing only battery voltage and circuit board temperature, the reorganizing system could institute random changes in the behavioral hierarchy (on the circuit board) whenever the sensed states of these variables departed from their reference states in either direction. This reorganization would cease only when the behaving system had learned to seek a perceptual state of affairs that, as a side-effect, provides enough sunlight to keep the battery optimally charged but not so much as to make the circuit board overheat.

Suppose that the environment contained signal lights: steady light means high temperature in a particular place, blinking light means a shadowed and cool place. But the control systems don't have to know about those meanings. Through reorganization, they would end up controlling for the blinking light being close some of the time, the steady light being close another proportion of the time, and other things the rest of the time. What the control systems end up controlling for has nothing _a priori_ to do with battery voltage or internal temperature. Neither the control systems nor the reorganizing system knows about the rule that says "steady means hot, blinking means cool and no light." The reorganizing system knows only about battery voltage and circuit temperature. The behaving systems know only about blinking and steady lights. The link between the lights and the internal variables is a secret of the physical environment, which neither the reorganizing nor the behaving systems have to discover. Ever.

The power of this arrangement is that the reorganizing system does not depend on the environment to contain any particular causal rules. If someone comes along and switches the lights, so that blinking means hot and steady means cool shade, the behaving system will initially seek the wrong places. But the effect will be lowering of the battery voltage or heating of the circuit board, and reorganization will start. It will end when the behaving system has switched which light it visits under what circumstances (or succumbs). The reorganizing system, set up as I visualize it, can survive under ANY ARBITRARY set of environmental rules connecting behavior to internal state, at least enough of the time to matter at the species level.

Having said all that, I should also say that reorganization will work better if there are sensory signals available to the behavioral hierarchy that indicate at least roughly some of the most critical intrinsic variables. It would be easier to design that robot for reorganizability if there were available a sensory signal entering the behavioral systems like any other sensory signal, one indicating battery voltage and another indicating circuit temperature. These sensory signals don't have to have any initial significance in the behavioral hierarchy. But if they're available, like "emotion" signals, the behavioral hierarchy can come to perceive relationships between them and exteroceptive signals, and it can learn to control for certain states of these relationships. Perceiving low battery voltage by the behavioral systems just means perceiving low battery voltage; that has no particular value to the behaving systems. That could be a good thing or a bad thing. The reorganizing system, however, which has not only its own perception of battery voltage but a built-in reference level for it, will cause reorganization of the behavioral systems until they learn to treat a low battery voltage signal as an error, and until they link this error to the adjustment of some sensed state of the external world that does, however indirectly, result in raising the battery voltage. After a while, this behavioral system will quickly sense low battery voltage and do what it has learned to do to increase that signal to its reference value, and it will do this before the reorganizing system sees enough error to cause reorganization to begin. That's why we learn to eat at mealtimes instead of when we're hungry. We've learned that if we eat regularly, a certain signal will be prevented from appearing. That signal is one which, if it gets large enough, results in an episode of reorganization and disorganization. So we say to ourselves that that is a "bad" signal. The feeling of hunger is a "bad" feeling.

There are lots of problems yet to be solved before a real test of my scheme can be devised. I'm plugging away with tests of various kinds of reorganizing in simple control systems just to find out what happens when you do it in different ways. Some ways are terribly slow; some work so fast that I'm suspicious of them. I've already ruled out one: reorganizing by switching the polarity of the output effects from 1 to minus 1 doesn't work, because it causes tremendous transients that can flip the system into unrecoverable states. Maybe if I stick a 0 in the middle as a third possibility it will work better. Or maybe output reorganization has to occur on a continuum between 1 and -1. We have to find these things out before we can come up with a workable system

My main point here is that reorganization has to be based on sensing variables with _a priori_ meaning to the organism. This rules out all, or essentially all, environmental variables -- all of the variables that the behavioral control systems are most concerned with controlling.

Best, Bill P.

Date: Thu Aug 20, 1992 9:40 am PST
Subject: Re: PCT & interactionism, Blindmen II

[From Oded Maler 920820]

(Rick Marken, ibid.)

Simon said that we have all kinds of great computer models of behavior -- so we understand behavior better than we think. My commentary said "that depends on what you mean by "behavior", Nobel man".

It depends also on what you mean by "understand" and by "think".

I'm just rereading a paper also written as a critique of similar claims by Simon - it's Hofstadter's "Waking up from the boolean dream" which I, btw, recommend. I think it won't hurt anybody if you re-post your paper. And I can't help reminding you that wrt to the universe/life, even the most knowledgable PCTer is also partly blind (partly by choice and partly because of difference in order of magnitudes between the observer and the observed).

I don't want by any means to interfere with your local politics, or to start a thread on these topics, but it seems to me that republicans are controlling by *all* the means they have for the obvious variable for a party - to be in power. If the effects of the crap they probably said in their convention works on the average (w. or w.o. brain) American cityzen is positive, why should their worry of the opinion and the aesthetic feelings of some marginal (quantitatively, of course) R.M. in L.A.?

Best regards --Oded

Date: Thu Aug 20, 1992 11:32 am PST
Subject: injury

(penni sibun920820.1100)

[Martin Taylor 920820 10:30]

Sorry about your repetitive motion injury. I do hope you are getting medical attention for it and not just trying to cure it by yourself (end parental mode).

yes; thanks. i'd like to point out that carpal tunnel syndrome is the fashionable disease of the 90s. however, most repetitive problems are *not* cts. unfortunately, most doctors, as well as lay people, tend to assimilate any problems with one's upper limbs to cts, which hinders rather than helps diagnosis and treatment. so if any of you have problems, please don't accept a cts diagnosis at face value: make sure it matches your symptoms and other conditions are ruled out.

--penni

Date: Thu Aug 20, 1992 11:48 am PST
Subject: categories

(penni sibun 920820.1200)

[From Mary Powers 9208.19]

So am I. Word/sentence/paragraph is as pseudo-hierarchical as cell/organ/organism/group. They are all at the category level.

i find this statement intriguing, and would be delighted if you would care to expand on it a bit.

--penni

Date: Thu Aug 20, 1992 4:31 pm PST
Subject: Re: Little Baby & Reorganization

[Martin Taylor 920820 19:15]
(Bill Powers 920818.2300)

Before your long tutorial on reorganization got here, Jeff, Allen, and I had a discussion on Jeff's proposal. I hadn't fully understood it from the written description, and I don't think you did, either. But your tutorial helped us to refine our understanding of your approach to reorganization. Let me see if I can rephrase Jeff's as I understand it. (If Jeff doesn't accept this as his, then consider it a derived version that I think should be considered).

I'm not saying I buy Jeff's system, but I don't find it incredible, either. It sounds workable, and has a good evolutionary rationale based on real long-term consistencies in the environment.

Jeff starts with the proposition that I use in my version of reorganization, that persistent and growing error in any ECS will cause THAT ECS to reorganize its local environment. There are all sorts of ways it could reorganize, none of which imply any knowledge about what it is doing within the hierarchy or in the outer world. We all agree on the need for that kind of ignorance (although I have, as we say, a "flag up" on that, for later discussion in connection with symbolic AI).

Jeff incorporates a separate control system, which could be a small hierarchy, that you call a "parasite." I'd prefer to call it a "symbiote." It is a very primitive thing, in evolutionary terms. It takes advantage of regularities that do occur in the environment, such as that too much thermal energy disrupts biologically active molecules, and that therefore it is a good idea for ANY organism to move away from sources of high-level thermal energy. The symbiote doesn't know how to make this move happen, so what it does is to execute (possibly random) actions in the motion effectors. It knows where they are, relative to the source of the unwanted energy, because it has evolved as part of the organism's ancestors for long enough that sucessful symbiotes have selectively survived. But it doesn't know any more than to create some strong action in,say, an arm or a leg. It is not controlling in the sense we usually use, with precision, but it is controlling in the sense that flailing around quite often moves one away from a localized source of trouble, and that brings the symbiote's percept into a range near its reference level. (Aside: control to be away from some central point is a quite different problem from control to be near the central point. It's not just a question of reversing the sign of the output. It's the difference between "No" and "Yes" in language. "No" means "Anything except what was proposed," a wide range of possibilities.)

What happens when the symbiote induces this relatively violent activity? ECSs experience large and momentarily growing error if they have as part of their feedback loop the actuators affected by the symbiote. This (in my view) is a condition that lasts for a moment, in which any affected ECS might experience a reorganization event. That might involve a change in their perceptual input function or in their output connections, or it might only affect their gains, or conceivably their reference input connections. But it would be a local change. These ECSs know nothing of the symbiote. But so long as occasions arise in which the symbiote is activated, they will reorganize, and thus will (if their perceptual input connections have the information) become less likely to act in such a way as to activate the symbiote.

The symbiotes are, as you point out, already set up to respond to particular environmental states, or, more properly, to internal states that in many cases have simple environmental correlates. All the symbiotes can do is to cause muscular activity in the world, or to cause the organism to emit chemicals or to change its appearance (the scent of fear, facial flushing or the colour change of a chameleon, for example). They do not, as you suggest, "teach another system to do it the same way." They have primitive stereotyped behaviour that is often immediately effective. But the evolved main hierarchy can learn much better ways of not getting into the bad situation, or of getting out of it once there.

Gordon Liddy can burn his hand, because he has a high-gain (high insistence) ECS for something of which that act forms a part. High-gain systems can over-ride the primitive symbiotes, but it is hard. Most people couldn't do it, except under extreme duress, where there are large deviations from many conflicting reference levels (excuse the shorthand).

You are quite right that an acting symbiote reduces the degrees of freedom available for the main hierarchy to act, possibly drastically. But that loss seems validated by experience. We do not maintain our flexibility of action in a "primitive stress" situation. We tend to focus on removing the stressful condition, which (with luck) the main hierarchy has learned to do. Adults don't kick and scream when hunger gets extreme. They try to find food, with some determination and ignoring other tempting distractions, at least until their systems lack the energy sources to keep going. Babies kick and scream, which, if Jeff is right, is the effect of the symbiotes. Adults focus on control that works in a much more complex environment than just the proximity of Mother.

The symbiote's job, as I see it, is not to signal "Pain" but to create error in the hierarchy, especially in a part whose reorganization is likely to result in useful control of whatever the symbiote does not like. "Pain" is an interpretation of stimulus conditions accessible to the main hierarchy, stimulus conditions that have been associated with action of the symbiote in earlier days (and perhaps still). There are many kinds of pain, and if Jeff's notion is right, maybe they are associated with different symbiotes.
========================

I think we have three candidate reorganization systems that are plausible, and I am happy to hear that this is what you are working on right now. We may actually be able to help with real (simulation) tests before too long (not me personally, but my faithful contractors).

The three, in cartoon form:

1) Powers: Intrinsic variables (such as body chemical states) have reference levels fixed by evolution. Errors in them induce a tendency to reorganize in the CNS hierarchy. The actions of the CNS hierarchy affect the levels of the intrinsic variables only as side-effects of whatever happens in the outer environment. There may possibly be some localization of the reorganization to

the upper levels of a growing hierarchy, leaving lower levels relatively immune. As the hierarchy grows, new, labile ECSs are added on top of existing ones, but layers can grow laterally as well.

2) Taylor: Intrinsic variables participate in a single hierarchy along with the CNS. The ECSs that have intrinsic variable levels as (part of) their perceptual input functions are always at the top level of a building hierarchy. Persistent and growing error in any ECS induces reorganization, which may include the generation of a new ECS that receives its reference signals from the reorganizing one and that gets its perceptual input and sends its reference output to places connected to the reorganizing ECS. Other aspects of reorganization include the making, breaking, and reweighting of output-to-lower-reference connections, and Hebbian modification of perceptual inputs.

3) Hunter: Intrinsic variables provide inputs to primitive symbiote "control" systems that have the main function of inducing error in ECSs in the main hierarchy when the intrinsic variables depart far from their reference levels. Reorganization in the main hierarchy proceeds roughly according to the Taylor model, except that the top-level ECSs are not connected with intrinsic variables, and new ECSs are "grown" on top of the hierarchy, in the same way as the Powers proposal. ======================== I'm plugging away with tests of various kinds of reorganizing in simple control systems just to find out what happens when you do it in different ways. Some ways are terribly slow; some work so fast that I'm suspicious of them. I've already ruled out one: reorganizing by switching the polarity of the output effects from 1 to minus 1 doesn't work, because it causes tremendous transients that can flip the system into unrecoverable states. Maybe if I stick a 0 in the middle as a third possibility it will work better. Or maybe output reorganization has to occur on a continuum between 1 and -1. We have to find these things out before we can come up with a workable system

The Little Baby is intended to provide a reorganization testbed, with a view to using its findings for developing much more complex networks for dynamic perceiving systems--speech recognition is the target. Chris Love is building it, not Jeff and Allan, but they do talk together. I wish you had a Macintosh, for 2 reasons: it is so much easier than a PC to work with, and we could communicate programs. We are using the Object-Oriented visual dataflow language Prograph. It's a bit hard to convert between that and a text-oriented language like C, though algorithmic concepts and formulae are relatively easy.

We would love to know anything you find out about reorganization from your experiments.

Maybe the transient when an output changes from -1 to 1 could be avoided if the gain dropped during the reorganization time for a period close to the inverse bandwidth of the feedback loop? Or perhaps the problem is that your hierarchy is too symbolic, in the sense that a given ECS has a function very distinct from all the others at its level, and therefore has only a very few connections at a lower level? If an ECS provided references to many lower ECSs that had highly correlated perceptual input functions, then the flipping of the sign of one link would not produce a large transient. It would be much like having a continuum of output weights, and I think it might also provide for more flexibility in the final (reorganized) system. But it would be computationally expensive to simulate.

Just a couple of suggestions.

Martin

Date: Thu Aug 20, 1992 4:57 pm PST
From: mmt
EMS: INTERNET / MCI ID: 376-5414
MBX: mmt@ben.dciem.dnd.ca

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Re: Mirror diagram

Dag,

I sent (or rather, asked our secretary to send) my Paris abstract last Tuesday (18th) let me know if you don't get it by, say the 26th (or if you do). All comments will be welcomed, but will be useless for the actual talk if received later than Sept 4. They may get incorporated in a paper that I hope the abstract will become, even if received later than that.

Martin

Date: Thu Aug 20, 1992 8:29 pm PST
Subject: language

(Penni Sibun 920820.2000)

> [Martin Taylor 920819 15:00]
>
> >i am deeply suspicious of hierarchical organization from my experience
> >in modelling language (as well as other behavior). language is
> >supposed to have all sorts of hierarchy in it (eg, a text is composed
> >of paragraphs are composed of sentences are composed of words) and i
> >just don't think this is a very useful way of looking at language, and
> >am working to show this.
>
> I would be very interested to know more. I have evolved in exactly the opposite direction. My theory of Layered Protocols seems to me to hold promise of clarifying many of the difficult problems of language, and I think that arguments from information theory and feedback stability theory almost demand that dialogue be conducted in a hierarchic way. If you have counterarguments, I'd be most interested in them.

i was talking about any text, written or spoken. are you talking about just dialog? i agree that dialog has properties that distinguish it from eg written text. but i quite disagree that dialog is hierarchic. (that's been the model in computational linguistics for nearly 20 years, and it has had approximately zero efficacy.) can you tell me what forms the structure of the hierarchy?

> But if all you mean is that one cannot describe language as the simple combination of unitary elements of different sizes, then I totally agree. Your use of the word "composed" suggests that this may be your meaning.

yes, i mean that--and more!

> All the elements, of whatever size, change their functions dramatically depending on situational and linguistic context, and the contextual effects can be very long-lasting.

i'm going to suggest that the elements aren't anything we might think. eg, in spoken language, ``word'' is not a useful concept.  in written language, it's only useful because ``word'' refers to elements of our encoding system--not of language itself. because we are literate people in a literate society, we tend to think of ``word'' as the stuff between spaces on a page, but these words don't particularly seem to be what we use when we utter stuff.

> Language is not a concatenative process.

i'm not exactly sure what you mean here, beyond a reaffirmation of hierarchy.

cheers.   --penni

Date: Thu Aug 20, 1992 8:46 pm PST

(penni sibun 920329.2200)

> [From Rick Marken (920820.0900)]
>
> Note-- I have re-written (slightly, but brilliantly) the "Blind men" paper as a commentary on an article by Herb Simon that appeared in a recent issue of Psychological Science. It is relevant to our current discussion of "interactionism". Simon said that we have all kinds of great computer models of behavior -- so we understand behavior better than we think. My commentary said "that depends on what you mean by "behavior", Nobel man". I really only added a couple paragraphs to the old "Blind man" paper. But if anyone is interested, I will be happy to post the new version. I've already submitted it to Psych Science but I'd still appreciate comments.

i'd be interested. it may not surprise you to know that simon has a (probably similar) paper coming up in _cognitive science_ and agre has written an excoriating reply.

cheers.   --penni

Date: Fri Aug 21, 1992 2:36 pm PST
Subject: Re: relation of effort to result

[From Bill Powers (920821.0800)]

Avery Andrews (920820) --

> As to `how many' regularities there are in nature, I doubt that this is a >question that can be profitably argued about. At a very minimum, PCT depends on >there being consistent correlations between the direction of gross effort and >the direction of change of net result, & a continuous function relating the >magnitudes of these two quantities.

We fall easily into treating nature as an abstract aspect of our models, as if it were so hard to find any natural phenomenon that we have to guess what nature would really look like if we could find some. You're right: we depend for control on consistencies in nature. I'd go farther -- we don't really expect natural laws to be statistical. If there are disturbances, we expect to find explanations for them and we expect that there will be some systematic way of opposing them. Disturbances don't happen on a time- scale or a scale of magnitude that prevents us from opposing most of their effects instead of just cleaning up after them. I suppose that defines our niche -- we perceive the world in such a way that we can handle its vagaries.

But all is not quite as obvious and regular as you imply. The correlation between gross effort and the direction of change of net result is not always as we assume it to be, nor is it the same from one circumstance to another. In BCP I brought up the example of a horse pulling a cart up a hill, crossing the top, and pulling it down the other side. When it's going down the other side, all its muscle forces are pushing in a direction opposite to their direction when going up the hill: the horse is moving forward by pushing backward. When you lower a bucket into a well, the forces produced by your muscles never push the bucket downward. When you slowly raise an outstretched arm from your side to the overhead position and slowly lower it again, the torques created by the muscles at the shoulder joint are always in the upward direction; when the arm is moving either up or down at a constant angular velocity, the torques are identical at a given position. When you stretch your arm out to point to something, the gross effect is in the direction of "reach" -- you can just FEEL that "reaching effort." But there is no vector force in the direction of pointing -- the net force is upward, counteracting gravity. The main sensation of reaching comes from trying to bend the elbow joint past its limit. There are no muscles in your arm that can cause it, when already outstretched, to lengthen, even though pointing "hard" feels that way.

When you swing that outstretched arm laterally, pointing first at one object and then at another, the sideward forces rise to accelerate the arm, drop to zero as the arm coasts toward the new direction supported upward torques at the joint, and then reverse to decelerate the arm to a stop -- all while the arm is swinging in the same direction.

When you use the steering wheel of a car to control a car's position on the road, there is no relationship between the steering wheel angle and the lateral position. The relationship is between wheel angle and the rate of change of sideward velocity -- the second time derivative of position. You might have to turn the wheel to the left in order to negotiate a right turn, if you're turning from a direction into the wind. The steering wheel turns in the direction of the curve at first. Then as the headwind becomes a crosswind, the angle of the wheel lessens more and more, until it actually opposes the turn (the crosswind alone is more than enough to make the car follow the turn of the road). The same thing can happen on entering a steeply banked turn when you're going around it at less than the design speed.

When Astro is making a final approach to Mother in more thabn one dimension, its jets are firing opposite to the direction of travel. When its path is curving and its velocity is changing, its jets could firing at any angle to the direction of travel.

Driving a screw into a piece of wood (after it is started) is done not by pushing it into the wood but by twisting it. Hammering a nail in requires, at the moment that the hammer head is about to strike the nail, pulling the handle toward you to counteract centrifugal force. The same is true of a golf swing -- just as the head of the club strikes the ball, you're pulling strongly on the grip to keep the head from continuing in a straight line into the ground. When you throw a baseball, you release it at the moment its path is toward the target, which is not when your arm is pointing at the target but when it is at a large angle to the direction to the target. The follow-through has no effect on the ball at all. When you go from a standing position to a squat, it is gravity that moves your body downward while your muscles exert efforts as if to stand up. When you press lightly downward on a tabletop with one finger, the main force in your arms is upward, supporting the arm and keeping the touch from becoming heavier. Just at the point where you're pressing down with a force equal to the weight of the arm times its average radius from the body, a force of several kilograms at the fingertip, the upper- arm muscles are mainly keeping the elbow from bending and the shoulder muscles are doing nothing. There is another force at the fingertip at which the shoulder muscles are supporting the upper arm and one end of the forearm, while the biceps and triceps are doing nothing but oppose each other.

The only way it can seem that gross effort corresponds to gross result is by confusing the effort with the result. We simply assume that if an effect is going in a given direction, the efforts we feel must be aimed in that direction, too. This is what I mean when I say we name behaviors after their controlled results, not after the efforts required to keep them controlled. We not only name them that way, but we interpret our experiences that way even without words. When you unlatch the door on a car that is parked so it's tilted toward you, you use the same term, "opening the door," for the action you take, even though you have to push on the door to keep it from opening too fast. We pay attention not to what we're doing, but to the effects of what we're doing. Because we control so well, we don't pay any attention to all the little disturbances that are present, even if counteracting their effects requires reversing our normal action to keep the same outcome going. Unless someone calls attention to them, the disturbances won't even be noticed. Because they're not noticed, the rather radical changes in action needed to cancel their effects are also not noticed. We end up with an unrealistic sense of how simple and straightforward behavior is. It's simple and straightforward only because we're excellent control systems.

I have on the page before me a diagram of a circuit that ought to work,
>and is supposedly based on the actual wiring of a real critter, and which I believe I described accurately.

If I showed you only the connections in an electronic circuit, without identifying the components or giving their values, neither you nor anyone else could say what that circuit does. It makes a great deal of difference, for example, whether the elements are flip-flops or continuous amplifiers. It makes a lot of difference whether a given synapse has 1 unit of effect on the following neuron or 1000 units, and whether the effect is proportional, rate-sensitive, or time- integrated. When closed loops can be traced out, it makes a lot of difference if the net sign around the loop is positive or negative; if positive, whether the gain is greater or less than 1, and if negative, whether it is only -1 or -2 or more like -100. When signals from different sources converge on a given neuron, it makes a difference whether their effects are both excitatory, or opposite in sign. Do we have an amplifying adder here, or a comparator, or a multiplier, or a gate? You can't see the function of any connection just by looking at its geometry.

None of these critical details can be derived simply from a drawing of the network of neural connections in Aplysia.

Best, Bill P.

Blindmen II

[From Rick Marken (920821.1800)]

In response to popular demand (3 people asked) I am posting the new version of the Blindmen paper (submitted to Psych Science, copyright me) -- it's a bit more than 20Kb (sorry Gary).

Here it is:

The Blind Men and the Elephant:
Three Perspectives on the Phenomenon of Control

Richard S. Marken
August 12, 1992

Abstract - The enthusiasm expressed by Simon (1992) about current explanations of behavior may be greater than is warranted by prevailing uncertainties about the nature of the phenomenon to be explained. Behavior has been described as a response to stimulation, an output controlled by reinforcement contingencies and an observable result of cognitive processes. It seems like these are descriptions of three different phenomena but they are actually descriptions of three different aspects of the same phenomenon -- control. Control is like the proverbial elephant studied by the three blind men; what one concludes about it, and how one tries to explain it, depends on where one stands. It is suggested that the best place to stand is where one has a view of the whole phenomenon - be it elephant or control.

In a recent article (Simon, 1992), Nobel Laureate Herbert Simon asked "what is an 'explanation' of behavior?" and answered with a tribute to the success of various computer-based, dynamic simulations of psychological processes. Simon's enthusiastic answer to his rhetorical question was most heartening but somewhat optimistic from the point of view of those, like myself, who believe that psychology has yet to answer a more fundamental question, namely, "what is behavior?". The tacit answer to that question seems to be "behavior is what organisms do". But it is possible to create circumstances where one can see every aspect of an organism's behavior and still be unable to tell what it is up to (Marken, 1989). Results like this suggest that successful simulation of the appearance of behavior can still fail to explain what an organism is doing. In order to understand why this might be the case it is necessary to examine the possibility that one answer to the question "what is behavior?" might be "behavior is what organisms control" (Powers, 1973).

The behavior of living organisms (and some artifacts) is characterized by the production of consistent results in an unpredictably changing environment, a phenomenon known as control (Marken, 1988). Control can be as simple as maintaining one's balance on uneven terrain or as complex as maintaining one's self-esteem in a dysfunctional family. Control is a pervasive aspect of all behavior yet it has gone virtually unnoticed in psychology. What has been noticed is that behavior appears to be a response to stimulation, an output controlled by reinforcement contingencies or an observable result of cognitive processes. Each of these appearances is what would be expected if people were looking at control from different perspectives. The situation is similar to that of the the three blind men who were asked to describe an elephant; the one near the tail described it as a snake, the one near the leg described it as a tree trunk and the one near the side described it as a wall. Each description gives an accurate picture of some aspects of the elephant, but a false picture of the elephant as a whole. If behavior involves control then psychology, too, has given an accurate picture of some aspects of behavior but a false picture of behavior as a whole.

Closed-Loop Control

The basic requirement for control is that an organism exist in a negative feedback situation with respect to its environment. A negative feedback situation exists when an organism's response to sensory input reduces the tendency of that input to elicit further responding. Negative feedback implies a closed-loop relationship between organism and environment; sensory input causes responding that influences the sensory cause of that responding, as shown in Figure 1. It is hard to imagine an organism that does not exist in such a closed-loop situation because all organisms are built in such a way that what they do affects what they sense. Eyes, for example, are located on heads that move so that what the eyes see depends on what the head does. To the extent that what the head does depends on what the eyes see (such as when the head turns in response to an attractive passer-by) there is a closed loop; sensory input causes responding (head movement) which affects the cause of responding (sensory input). The feedback in this loop must be negative because behavior is typically stable (organisms, for example, do not normally exhibit the "run away" behavior that characterizes positive feedback loops, such as the "feedback" from a microphone that amplifies its own output).

Insert Figure 1 About Here

The fact that organisms exist in a closed negative feedback loop means that two simultaneous equations are needed to describe their relationship to the environment. These are given as equation (1) and equation (2), below. The terms in these equations are summarized for reference in the discussion that follows:

s = sensory input variable
r = response variable
s^* = reference value for sensory variable such that $r = 0$ when $s = s^*$
d = environmental variable
$k.o$ = organism function relating sensory variable, s, to response variable, r
$k.e$ = environmental function relating environmental variable, d, to sensory variable, s
$k.f$ = feedback function relating response variable, r, to sensory variable, s

For simplicity we will assume that all functions are linear and that all variables are measured in the same units.

Equation (1) describes the effect of sensory input on responding so that:

$$r = k.o\,(s^*-s) \tag{1}$$

This equation says that responding, r, is a linear function of sensory input, s. The sensory input is expressed as a deviation from the value of input, s*, that produces no responding; s* defines the zero point of the sensory input. Equation (2) describes the effect of responding on sensory input. For simplicity it is assumed that responding, r, adds to the effect of the environment, d, so that:

$$s = k.f\,(r)+ k.e\,(d) \tag{2}$$

The variables r and d have independent (additive) effects on the sensory input, s. The nature of the environmental effect on sensory input is determined by the environmental function, k.e. The feedback effect of responding on the sensory cause of that responding is determined by the feedback function, k.f.

Equations (1) and (2) must be solved as a simultaneous pair in order to determine the relationship between stimulus and response variables in the closed loop (the derivation is shown in the Appendix). The result is:

$$r = 1/((1/k.o)+k.f)\ s^* - k.e/((1/k.o)+ k.f)\ d \tag{3}$$

Equation (3) can be simplified by noting that the organism function, k.o, transforms a small amount of sensory energy into a huge amount of response energy (such as when a pattern of light on the retina is transformed into the forces that move the head). In control engineering, k.o is called the "system amplification factor" or "gain" and it can be quite a large number. With sufficient amplification (such that k.o approaches infinity) the (1/k.o) terms in equation (3) approach zero, so equation (3) reduces to:

$$r = s^*/k.f - (k.e/k.f)\ d \tag{4}$$

Equation (4) is an input-output equation that describes the relationship between environmental (stimulus) and response variables when an organism is in a closed- loop, negative feedback situation with respect to its environment. The result of being in such a situation is that the organism acts to keep its sensory input equal to s*, which is called the reference value of the input. Equation (4) shows that the organism does this by varying responses, r, to compensate for variations in the environment, d, that would tend to move sensory input away from the reference value; this process is called control.

Three Views of Control

All variables in equation (4), with the possible exception of s*, are readily observable when an organism is engaged in the process of control. The environmental variable, d, is seen as a stimulus, such as a light or sound. The response variable, r, is any measurable result of an organism's actions, such as bar pressing or speaking. The reference value for sensory input, s*, is difficult to detect because an observer cannot see what an organism is sensing. But s* is the central feature of control since everything an organism does is aimed at keeping its sensory inputs at reference values. Because these reference values are difficult to detect it will not be obvious to an observer that an organism is engaged in the process of control. What will be obvious is that certain variables, particularly the environmental and response variables and the relationship between them, will behave as described by equation (4). Thus, equation (4) can be used to show what control might look like if one did not know that it was occurring. It turns out that there are three clearly different ways of looking at control depending on which aspect of the behavior described by equation (4) one attends to.

1. The stimulus - response view. This view of control sees behavior as a direct or indirect result of input stimulation. An example of stimulus-response behavior is the so-called "pupillary reflex" where changes in a stimulus variable (illumination level) lead to changes in a response variable (pupil size). The stimulus-response view is the basis of several current approaches to understanding behavior, such as the "synergistic" or "coordinative structure" theory of motor coordination. Warren, Young and Lee (1986), for example, describe a synergistic model of running in which "vertical impulse is directly modulated by the optical variable Ft..." (p.264). The behavior of running is seen in stimulus- response terms; a stimulus variable, Ft, determines ("modulates") the value of a response variable, vertical impulse. The stimulus-response view is also the basis of a recent theory of attention (Cohen, Dunbar and McClelland, 1991) in which connections between printed word stimuli and verbal responses in the Stroop effect are modulated by connections in a neural network.

Equation (4) shows that behavior will look like a stimulus-response process when the reference value for sensory input, s*, is a constant; for simplicity assume that it is zero. Then responding is related to environmental stimuli as follows:

(5) $r = - (k.e/k.f)\ d$

Equation (5) shows that, when there is a fixed reference level for sensory input, it will look to an observer of behavior as though variations in an environmental stimulus, d, cause variations in a response, r. This is what we see in the pupillary reflex where pupil size, r, is proportional to illumination level, d. Of course, this relationship between pupil size and illumination level is precisely what is required to keep a sensory variable (sensed illumination) at a fixed reference value (s* = constant).

One's inclination when looking at an apparent relationship between stimulus and response is to assume that the nature of that relationship depends on characteristics of the organism. Equation (5) shows, however, that when an organism is engaged in control, this relationship depends only on characteristics of the environment (the functions k.e and k.f); the organism function, k.o, that relates sensory input to response output, is rendered completely invisible by the negative feedback loop. This characteristic of the process of control has been called the "behavioral illusion" (Powers, 1978).

2. The reinforcement view. This view of control sees behavior as an output that is shaped by contingencies of reinforcement. A reinforcement contingency is a rule that relates outputs (like bar presses) to inputs (reinforcements); in equation (4) this contingency is represented by the feedback function, k.f, that relates responses to sensory inputs. The reinforcement view is the basis of at least one influential theory of generalization and discrimination (Shepard, 1987). In a connectionist implementation of the theory, a reinforcement contingency is used to shape the formation of generalization gradients (Shepard, 1990). The reinforcement view is also the basis of modern theories of operant behavior. According to Domjan (1987) the contemporary perspective on operant behavior focuses on how contingencies "restrict freedom of action and ... create redistributions of various types of activities"(p. 562). In other words, contingencies shape (redistribute) responses (activities).

Equation (4) shows that it will look like contingencies (the feedback function) control responses when s*, d and k.e are constants, as they are in the typical operant conditioning experiment. In these experiments, s* is the organism's reference value for the sensory effects of the reinforcement. The environmental variable, d, is the reinforcement, which, if it is food, is typically a constant size and weight. The sensory effect of a reinforcement can be assumed to be directly proportional to its size and weight, making k.e = 1. So, for the operant conditioning experiment, equation (4) can be re-written as

(6) $$r = S^*/k.f - D/k.f$$

where S* is the constant reference value for sensed reinforcement and D is the constant value of the reinforcement itself.

The only variable in equation (6) is the feedback function, k.f, which defines the contingencies of reinforcement. One simple contingency is called the "ratio schedule" in which the organism receives a reinforcement only after a certain number of responses. The ratio corresponds to the function k.f in equation (6). When the ratio is not too demanding it is found that increases in the ratio lead to increased responding. More demanding ratios produce the opposite result; increases in the ratio lead to decreased responding (Staddon, 1979). Either of these results can be produced by manipulating the relative values of S* and D in equation (6). The important point, however, is that the apparent dependence of responding on the feedback function, k.f, is predicted by equation (6). To an observer, it will look like behavior (responding) is controlled by contingencies of reinforcement. In fact, the relationship between behavior and reinforcement contingencies exists because the organism is controlling sensed reinforcement; responding varies appropriately to compensate for changes in the reinforcement contingency so that sensed reinforcement is kept at a constant reference value, S*.

3. The cognitive view. This view of control sees behavior as a reflection or result of mental plans or programs. This kind of behavior is seen when people produce complex responses (such as spoken sentences, clever chess moves or canny investment decisions) apparently spontaneously; there is often no visible stimulus or reinforcement contingency that can be seen as the cause of this behavior. The cognitive view is the basis of numerous psychological theories that propose mental algorithms to explain the appearance of cognitive behavior. Examples of such theories include the ACT (Anderson, 1983) and SOAR (Newell, 1990) models of cognition, hierarchical models of the generation of movement sequences (Rosenbaum, Kerry and Derr, 1983), connectionist models of speech production (Jordan, 1989) and schema models of expertise in problem solving (Lesgold, A., Robinson, H., Feltovitch, P., Glaser, R., Klopfer, D. and Wang, Y., 1988).

Cognitive behavior is most obvious when environmental factors (such as stimulus variables and environmental and feedback functions) are held constant. When this is the case, equation (4) becomes

$$r = s^*/F + K \qquad (7)$$

where F is the constant feedback function and K = (k.e/k.f) d, a constant. Since s* is typically invisible, equation (7) shows that there will appear to be no obvious environmental correlate of cognitive behavior. An observer is likely to conclude that variations in r are the result of mental processes -- and, indeed, they are. But it is actually variations in s*, not r, that are caused by these processes; variations in r being the means used to get sensory inputs equal to s*. Thus, chess moves are made to keep some sensed aspect of the game at its reference value. When the environment is constant, r (the moves) may be a fair reflection of changes in the reference value for sensory input. However, under normal circumstances r is only indirectly related to s*, variations in r being mainly used to compensate for variations in the environment that would tend to move sensory input from the reference value, s*.

Looking at the Whole Elephant

The blind men never got a chance to see the whole elephant but if they had they would have instantly understood why it seemed like a snake to one, a tree trunk to another and a wall to the third. Psychologists, however, can take a look at control and see why the appearance of behavior differs depending on one's perspective. What is common to the three views of behavior discussed in this paper is the reference for the value of sensory input, s*. Organisms behave in order to keep sensory inputs at these reference values (Powers, 1989). They respond to stimulation in order to keep the sensory consequences of this stimulation from moving away from the reference value; so it appears that stimuli cause responses. They adjust to changes in reinforcement contingencies by responding as needed in order to keep the sensory consequences of reinforcement at the reference value; so it appears that contingencies control responding. And they change their responding in order to make sensory input track a changing reference value for that input; so it appears that responding is spontaneous.

What appear to be three very different ways of describing behavior can now be seen as legitimate ways of describing different aspects of one phenomenon -- control. Each is just a different way of describing what an organism must do to keep its sensory inputs at their reference values. Indeed, once you know that the appearances called "behavior" are merely the visible consequences of an organism's efforts to control it sensory inputs, the problem of explaining behavior changes completely, from an attempt to build models that simulate the appearance of behavior (S- R, reinforcement or cognitive) to an attempt to build models that control the same sensory inputs as those controlled by real organisms. In order to build the latter type of model it is necessary to learn what sensory variables are actually being controlled by organisms. This type of investigation cannot be done by simply looking at the appearance of behavior. Methods based on control theory can be used to test which sensory variables an organism might be controlling at any time (Marken, 1992). These methods make it possible to take off the blindfolds and see the whole elephant -- the phenomenon of control.

Appendix

Given the two system equations:

(1) r = k.o (s*-s) and

(2) s = k.f (r)+ k.e (d)

we want to solve for r as a function of s. First, substitute equation (2) for s in equation (1) to get:

(A.1) r = k.o (s*-(k.f (r)+ k.e (d)))

Multiply through by k.o to get:

(A.2) r = k.o (s*) - k.o k.f (r) - k.o k.e (d)

Move all terms with r to the left side of the equation to get:

(A.3) r + k.o k.f (r) = k.o (s*) - k.o k.e (d)

Factor r out of the left side of the equation to get:

(A.4) r (1 + k.o k.f) = k.o (s*) - k.o k.e (d)

Divide both sides of the equation by (1 + k.o k.f) to get:

(A.5) r = k.o/ (1 + k.o k.f) s* - k.o k.e/(1 + k.o k.f) d

Finally, divide k.o out of the numerators on the right side of (A.5) to get equation (3):

(3) r = 1/((1/k.o)+k.f) s* - k.e/((1/k.o)+ k.f) d

References

Anderson, J. R. (1983) The architecture of cognition. Cambridge, MA: Harvard University Press

Cohen, J. D., Dunbar, K. and McClelland, J. L. (1991) On the control of automatic processes: A parallel distributed processing account of the Stroop effect. Psychological Review, 97, 332-361

Domjan, M. (1987) Animal learning comes of age. American Psychologist, 42, 556 - 564

Lesgold, A., Robinson, H., Feltovitch, P., Glaser, R., Klopfer, D. and Wang, Y. (1988) Expertise in a complex skill: Diagnosing X-ray pictures. In M.T.H. Chi, R. Glaser and M.J. Farr (Eds.) The nature of expertise. Hillsdale, NJ: Erlbaum

Marken, R. S. (1988) The nature of behavior: Control as fact and theory. Behavioral Science, 33, 196- 206.

Marken, R. S. (1989) Behavior in the first degree. In W. Hershberger (Ed.) Volitional action: Conation and control. Amsterdam: North Holland (pp. 299-314)

Marken, R. S. (1992) Mind Readings: Experimental studies of purpose. Gravel Switch, KY: CSG Publishing

Newell, K. (1990) Unified theories of cognition. Cambridge, MA: Harvard University Press

Powers, W. T. (1973) Behavior: The control of perception. Chicago: Aldine

Powers, W. T. (1978) Quantitative analysis of purposive systems: Some spadework at the foundations of scientific psychology. Psychological Review, 85, 417-435

Powers, W. T. (1989) Living control systems. Gravel Switch, KY: CSG Publishing

Rosenbaum, D. A., Kerry, S. and Derr, M. A. (1983) Hierarchical control of rapid movement sequences. Journal of Experimental Psychology: Human Perception and Performance, 9, 86 - 102

Jordan, M. I. (1989) Serial order: A parallel, distributed processing approach. In J. L. Elman and D. E. Rumelhart (Eds) Advances in connectionist theory: Speech. Hillsdale, NJ: Erlbaum

Shepard, R. N. (1990) Neural nets for generalization and classification: Comment on Staddon and Reid (1990). Psychological Review, 97, 579 - 580

Shepard, R. N. (1987) Toward a universal law of generalization for psychological science. Science, 237, 1317 - 1324

Simon, H. A. (1992) What is an "explanation" of behavior? Psychological Science, 3, 150 - 161

Staddon, J. E. R. (1979) Operant behavior as adaptation to constraint. Journal of Experimental Psychology: General, 108, 48-67

Warren, W. H. Jr., Young, D. S. and Lee, D. N. (1986) Visual control of step length during running over irregular terrain. Journal of Experimental Psychology: Human Perception and Performance, 12, 259 - 266

Figure Caption

Figure 1. Closed-loop feedback relationship between an organism, represented by the rectangle, and its environment, represented by the arrows outside of the rectangle. A sensory variable, s, influences responding, r, via the organism function, k.o,. Responding influences the sensory variable via the feedback function, k.f. The sensory variable is also influenced by an environmental variable, d, via the environmental function, k.e.

Figure 1

Date: Fri Aug 21, 1992 9:04 pm PST
Subject: agre, facism

[From Rick Marken (920821.2100)]

penni (and anyone else):

I get the impression that there is some book or article by Agre from which I can learn the principles of interactionism. Could you tell me what that is? I'll be at ucla tommorrow so maybe I can take a look at it. Maybe there's a forth blind man (the zen master, perhaps) that I neglected in my paper. It just struck me that the three views of behavior described in my paper probably don't include the interactionist view. It might be hard to represent that view in my scheme since it seems that interactionists make no distinction between organism and environment -- thus eliminating both equations 1 and 2. That would sure take care of my thesis.

By the way. Didn't anyone else feel the sense of relief that I felt after the republican convention when they revealed that all the problems in this country are the result of homosexuals and atheists -- and NOT JEWS. I'm particularly relieved that there has been no embarassing outcry agaist the republicans -- as there was agaist Joe McCarthy. Could the republicans have finally found the people that everyone can hate. I'm impressed; they are already AHEAD of Hitler -- and after only 48 years. Heil (er... Viva) Bush.

Best regards   Rick

Date: Sat Aug 22, 1992 2:27 am PST

I've been thinking a bit about steering, & here's what I've come up with so far. The only direct and simple relationship is between the absolute angular position of the wheel, and the curvature of the car's path, but I think that people pay little, if any, attention to either of these factos. What probably does matter are such things as:

a) the difference between where you are and where you want to be (transverse to the road), P*

b) the rate at the car is drifting across the road, dP*.
There may be serveral ways of picking this up (such as differentiating P*, or watching the motions of lines on the road w.r.t. the car hood, & I'm sure other things as well).

c) the rotational velocity of the steering wheel.

d) the angle between the orientation of the car and the tangent to the road.

e) how far the wheel has turned since some time t0.

I'll only consider (a-c), tho (d-e) seem like they might play a role too.

Observe that on a big lane change you don't normally try to get into your new lane as quickly as possible, but drift transverse to the road at a fairly steady rate. And also, that when you get close to where you want to be, the drift rate needs to be smoothly decreased to zero. So I propose that P* is controlled by means of controlling dP*. So dP*_ref should be a function of P*, zero when P* is zero, small when P* is small, but flattening off to a constant fairly quickly under normal circumstances (I won't speculate about panic swerves).

Then dP* itself gets controlled by controlling the perceived rate of motion of the steering wheel: when dP* is too big in the rightward direction, crank the wheel counterclockwise, & vice versa (I suspect that there are various ways in which the reference rate of wheel-rotation might be determined, and that it might even be possible to devise experiments to figure out which one(s) people were actuallay using). Lower-level systems take care of summoning up enough torque to get the wheel to rotate at the right speed (and producing the hand-over-hand motion needed to effect a large total rotation).

Hence there is a three-level hierarchy of controlled perceptions: relative position, rate-of-change of position, and rate-of-rotation of the steering wheel. And of course there'd be some more levels involved in producing the wheel movements. Monitoring the rate of change of (transverse) position makes it unnecessary to do the otherwise complex calculations that would be needed to figure out how fast to turn the steering wheel when, and monitoring the rate of rotation of the steering while makes it unnecessary to predict how much force needs to be applied to the wheel to get the required curvature of the path.

A further elaboration: it seems plausible that drivers also assess the curvature of the upcoming road, and anticipate how much turning of the steering wheel will be needed to keep the car on course (it ought to be possible to do a video-game experiment to see whether they actually do this). So there could be an anticipator circuit assessing the appearance of the road ahead, and producing reference levels for rate of wheel-turning. The higher-level feedback loop can then add an additional term to correct for the anticipator's errors. Having control for a rate-of-turn of the steering wheel means that this anticipator can be a lot simpler than it would otherwise have to be, since it doesn't have to guess how much force ought to be applied to the wheel (though there could be anticipators contributing the force reference levels, they wouldn't have to be very precise).

Avery.Andrews@anu.edu.au

Date: Sat Aug 22, 1992 7:43 am PST
Subject: Re: Politics and cars

[From Bill Powers (920822.0800)] Rick Marken (920821.2100) --

"Republicans" is, of course, a category, as is "Democrats" and "Jews." According to TV reporters, many people who claim to be instances of the Republican category were angry at Pat Buchanan's speech and embarrassed at much of the rest of the proceedings. I expect there were some Democrats who were embarrassed at the proceedings at the other convention. As with most perceptions based on class membership and other superficial "population sampling" methods and labels, it's not likely that anything you say about either Democrats or Republicans is true of any individual Democrat or Republican you happen to meet, or any Jew, either.

System concepts tend to be unarticulated and largely unaware. I suspect that everybody believes in "family values," even the Mafia. But it's hard to say what a family value is, and more to the point, what's good about it. What most people mean, I suspect, is "people like us." As Bush said, people are raised to believe as their parents believed, and as their parents before them. Apparently, to many people, being raised to believe in something is a better reason to believe it than reasons grounded in knowledge or understanding. The difficulty is that not everyone is raised to believe in the same things, so to argue for a belief on the grounds that one was raised that way is to admit equal justification for contrary beliefs held by people raised a different way. One would think that this principle would lead to great tolerance of others and the realization that a belief is, after all, only a belief, not knowledge. Apparently, however, it leads in the opposite direction. The assumption is often that the way I was raised is the right way, and everyone else is misguided, perverted, ignorant, or evil.

This analysis cuts both ways in any politico-social confrontation. As long as people act out of ill-formed and self-contradictory system concepts, they will not understand what is good about some principles and bad about others, save for the way one happens, by accident of birth, to feel about them. People will simply be vehicles in which other people's ideas from long ago replay themselves like recordings -- and, of course, they are recordings, reference signals being, I think, derived in large part from remembered experiences. At the highest level there's not much room for other ways of picking them.

Politics is fertile ground for a control theorist. What's needed is a study in depth of many individuals who call themselves by some party label. What are they controlling for? What principles do they uphold, and what methods do these principles justify? What kinds of errors do they perceive in various social situations, and how do they see their proposed actions as correcting those errors? What do they think of as the good life? As being a good person? How do they think people work, with respect to rewards and punishments, self-interest, social interest, and so on? Each person's structure of perceptions and goals at the higher levels explains how that person acts and what that person strives for. A study that investigates these structures in many individuals of different avowed political views could come up with many of the reasons for our social difficulties, in the form of contradictory goals, misperceptions, and methods of control that don't actually work. Once the real reasons for failure of our social systems were brought out, perhaps the way to a better solution to our problems would become more apparent. It isn't that social problems are so difficult. It's that our approach to them is confused and self- defeating.

Avery Andrews (920822) --

Making a computer model of a car for use by proposed steering control systems shouldn't be too difficult. The steering wheel input to the car can be make fairly realistic. A torque applied to the steering wheel will cause the wheel to rotate until the restoring force equals the applied torque. The restoring force will be proportional to the centrifugal force, which is m*r^2/v, r being the radius of the turn (computed from the steering wheel angle) and v being the tangential velocity of the car. The radius of the turn is computed from the distance between front and rear wheels and the angle of the front wheels: where lines perpendicular to the rear and front wheels intersect is the momentary center of curvature. This ignores factors like tire slip and so on, but will be plenty close for the purpose. You could probably just say that the radius of curvature is inversely proportional to the steering wheel angle and be plenty close. Empirically adjusted constants will make the behavior of the car realistic enough.

At a low level, the human control system can sense and control either the force applied to the steering wheel, or the steering wheel angle. Different control properties will result, and I'm sure do result in different drivers who solve the problem in different ways. Controlling applied torque will result in a wheel angle that depends on speed, whereas controlling wheel angle directly, rather than torque, will make the wheel angle independent of speed. I think the former will probably work better. A kinesthetic variable may prove important: the feeling of lateral acceleration that you get in a turn.

The rate of drift of the car sideways is, as you say, obtainable from dp/dt. Unless you want to make a whole visual model, the easiest way to handle this is just to say that there's a perceptual signal equal to dp/dt and not worry about how it's derived.

>Hence there is a three-level hierarchy of controlled perceptions: >relative position, rate-of-change of position, and rate-of-rotation of >the steering wheel.

Sounds workable to me. The easiest way to get the reference rate of rotation of the steering wheel is to integrate the output of the rate-of-change-of-position system. In other words, make the steering wheel system control the angle of the wheel, which is easy, rather than the rate of change of angle. By integrating the output of the higher system, you'll get a reference signal that keeps changing as long as there's an error in rate of turning -- which is just what you want, because a fixed wheel angle corresponds to a particular rate of turning of the car. When the rate of turning error is zero, the integral of the output will be constant, holding the wheel at a constant angle.

I think there's one missing step in here. If there's a position error (relationship level), the car has to be moved sideways by a fixed amount, ending up going in the same direction as before. To get there, its direction has to be changed. Any change in direction will cause sideward drift rate. So the next level down should control drift rate, not direction. The drift rate is varied by varying the car's angle, so car angle should be the next lower variable. And car angle is varied by varying rate of change of car angle, which is varied by controlling steering wheel angle. You'll end up with two derivative-perceiving systems, because position involves two time-integrals of steering wheel angle.

Of course the system could be designed more compactly with derivative and integral calculations inside one system, but the point is to get an arrangement that works before finding one that works the right way.

With regard to assessing curvatures of the road ahead, I finally figured out the simplest way to anticipate curves. It's to perceive the position of the car's hood relative to the road, but not the road right under the hood of the car. If you simply look at the road 30 or 40 feet ahead of the car, at a place in line with the body of the car, the curvature will be perceived automatically a little in advance of the time when the wheel has to begin turning. When you're going faster, you look farther ahead. From experience you learn how far ahead to set your perception point at various speeds. This eliminates the need for complicated perceptual calculations.

Have fun. Best to all, Bill P.

Date: Sat Aug 22, 1992 9:09 am PST
Subject: politics and PCT

[From Rick Marken (920822.1000)]

Why do I feel like Peter being patiently corrected (as usual) by Jesus? I spew a little sarcastic spleen and my mentor gently points out that I am tilting at categories -- just like the targets of my tilt.

But my little tilt might have served a purpose, viz., re-starting a dialog on (and, possibly, some exporation of) those "higher levels" of control that we all know and love.

Bill Powers (920822.0800) says:

>One would think that this principle
>would lead to great tolerance of others and the realization that a
>belief is, after all, only a belief, not knowledge. Apparently,
>however, it leads in the opposite direction. The assumption is often
>that the way I was raised is the right way, and everyone else is
>misguided, perverted, ignorant, or evil.

Excellent point. Why, indeed, might this be the case -- from a PCT perspective? How might one be able to help a person see this -- before he or she dispatchs you to allah, ships you to a concentration camp, fires you from your job or votes for people who will push laws that keep you out of the community?

>Politics is fertile ground for a control theorist. What's needed is a
>study in depth of many individuals who call themselves by some party

>label. What are they controlling for? What principles do they uphold,
>and what methods do these principles justify? What kinds of errors do
>they perceive in various social situations, and how do they see their
>proposed actions as correcting those errors? What do they think of as
>the good life? As being a good person? How do they think people work,
>with respect to rewards and punishments, self-interest, social
>interest, and so on?

Maybe we could get a volunteer "republican" from the audience? I don't call myself by any party label but I would be willing to by analyzed as a member of the PCT party -- indeed, I think there are what would be called "politcal" implications of PCT; maybe that's what I'm trying to get us to discuss -- the socio- political implications of the PCT model of human nature.

>Once the real reasons for failure of our social systems
>were brought out, perhaps the way to a better solution to our problems
>would become more apparent. It isn't that social problems are so
>difficult. It's that our approach to them is confused and self-
>defeating.

Yes, the solution to our social problems are simple -- just like the solution to our personal problems (internal conflicts). It's just a slight change in the rate of firing of a reference neuron or two. But, as people who are in conflict know (at least, symtom- atically), that little change is very difficult to make when you try to change it at the level of the conflict itself; control systems are VERY tenacious.

Best regards Rick

Date: Sat Aug 22, 1992 1:11 pm PST
Subject: Re: agre, facism

(penni sibun 920822)

I get the impression that there is some book or article by Agre from which I can learn the principles of interactionism. Could you tell me what that is? I'll be at ucla tommorrow so maybe I can take a look at it.

the only published book is

david chapman, _vision, instruction, and action_, mitp, 1991.

it's about sonja. doesn't go into a lot of the philosophy, but if you spent .5hr w/ it, you could get some of the back ground.

agre's thesis and the pengi papers are what i recommend. avery posted how to get the former.

i really recommend preston's paper. it starts at exactly the same place as yr blind men paper, adn draws a v. diff conclusion.

It just struck me that the three views of behavior described in my paper probably don't include the interactionist view. It

   might be hard to represent that view in my scheme since it
   seems that interactionists make no distinction between organism and
   environment -- thus eliminating both equations 1 and 2. That would

there is no interactionist manifesto, or even well-defined canon. (there's pockets of interactionist work in phil, psych, anthro, ling, ai....) i'll also remind you that i've presented the most radical view of interactionism possible. a risky pedagogical move--though if, after having thoroughly jumped down my throat, you go off to investigate on your own....

cheers. --penni

Date: Sat Aug 22, 1992 1:27 pm PST
Subject: Preston reference?

[From Rick Marken (920822.1400)]

penni --

I'd love to see the Preston paper -- especially if it's related to the thesis of the Blind men paper. Do you have a detailed reference? Also, I couldn't find Chapman's book at ucla so it looks like I'll have to get all these references (including Agre's thesis) through the library at work. But I'm really interested in Preston. Who is he/she. Where's the paper/book?

Thanks Rick

Date: Sat Aug 22, 1992 7:23 pm PST
Subject: Re: Preston reference?

(ps 920822.17)

beth preston is at the univ of ga, w/ a phil/ai appt. send her email asking for a preprint of the "3rd alternative" paper. i don't guarantee her sending it, but ask nice and drop my name. (she can prob. send it via email.)

    epreston@aisun3.ai.uga.edu --penni

Date: Sun Aug 23, 1992 7:16 am PST
Subject: Hierarchies

[from Mary Powers 9208.22]
(penni sibun)

The "hierarchy" words, sentences, paragraphs:

>A sentence is a collection of words, a paragraph a collection of sentences.

A word can be a sentence or a paragraph, a sentence a word or a paragraph, a paragraph a word or a sentence.

Cell, organ, organism, group:

An organ is a collection of cells, an organism a collection of organs, a group a collection of organisms. At least in the opinion of some people.

A cell can be an organism, or vv. However an organ is not viable on its own, and (most?) PCTers don't buy groups as superorganisms. But whoever made this list, and the many who have gone along with it, probably look at it like word-sentence- paragraph.

The problem in each case is thinking of larger and larger collections as hierarchies. What Bill has tried to do in HPCT is develop a hierarchy of levels of perception in which the level below is necessary for the level above, but the higher level is of a different type. You must have sensations in order to perceive a configuration, but a configuration isn't simply a collection of sensations.

I'm not really equipped to get further into this. However, by a happy coincidence, Bill was writing a commentary to Martin Taylor which covers this very subject, and I've asked him to put it on the net here:

[Slightly edited ...WTP]

You'll remember that some time ago I raised the point that not all levels of the hierarchy would work with weighted sums of lower-level signals. The logical levels obviously don't. This has a bearing on your discussion of conflict, cooperation, and coordination.

The levels as I have defined them have a characteristic that's not immediately obvious. There is no generating principle which, given the definition of one level, would predict the definition of the next. These levels were found empirically, by a slow process of observing myself perceiving and doing things that weren't in the model yet. Each level came as a surprise and even a shock as I realized how obvious was the type of perception that I'd been taking for granted.

I think that my levels are orthogonal to the levels below them. This means that each new level introduces new degrees of freedom. If a new level used an input computation of the same type as that of the lower level, this would not be true. But when a new type of perception is brought in, the same algebraic considerations cease to hold.

Consider sensations versus configurations. A sensation is a weighted sum of intensity signals. So the number of possible sensations is the number of different and independent weighted sums. If there are nine intensity signals, there can be nine independent sensation signals at a given time. Let's suppose that these nine sensations are discriminable tones of a scale. Clearly, only nine notes at a time can behave as independent tones.

A configuration in the auditory modality is a static pattern of sensations. For example, a major triad chord would qualify as an auditory configuration. With nine notes available, we could have three simultaneous independent triad chords, sometimes overlapping some notes and sometimes not. But these triads are not weighted sums; they are weighted products.

If the sensations designate frequencies, a chord is defined by f2 = K1*f1 and f3 = k2*f2. Actually, it's probably more like f2 = (K1 mod 2)*f1 and so on, because of octaves. So the rules of linear combination no longer hold true. It's possible for a given set of nine configuration-perceivers to respond to different subsets of nine sensation-signals depending on the frequency range in question. When you introduce higher powers of the variables, the number of degrees of freedom multiplies (no pun intended).

Some of these degrees of freedom consist of variations at the sensation level that make no difference at the configuration level, but nontheless are present in it, or pass through it.

If the next level is transitions, or paths, we now have what seem like infinite possibilities in the ten configuration signals. We have beat notes, glissandos of whole chords or single notes, vibrato, tremolo, accelerando, diminuendo, and other variations that mean nothing at the configuration level. We now have combinations of the various derivatives which are not sensed at the configuration level. Each derivative adds a new degree of freedom.

And if the next level is events, we have phrases that include all the possible combinations of the variables at the transition, configuration, sensation, and intensity levels -- which looks like a very great deal more than nine degrees of freedom altogether. We have all kinds of frequency-time patterns, and at each level something new is introduced that is not constrained at the lower levels.

So I think that the degrees-of-freedom argument based on repeated application of weighted summations to previous linear combinations of weighted summations does not apply to the levels of perception as I have defined them. This is because with each new level, as I have defined levels, there is a change of logical type; no new level of perception is commensurate with any lower level.

At each new level, a new consideration is introduced: the most obvious example is the sequence level. What are the possible dimensions of ordering? I think that they are probably of a different order of infinity from the dimensions of lower-level perceptions. The number of possible orderings of three triads is limitless, not even considering possible orderings of possible time-frequency patterns. This has become a combinatorial problem, not an algebraic one: there are as many dimensions are there are elements to be ordered. Any set of configurations can represent a snapshot of an infinity of different sequential progressions. And when the logic level is introduced, we have 2^(2^n) possible Boolean expressions with n elements; elements connected by "and", "or", and "not" can't be expressed in the same terms as Ax + By + Cz. And that only takes care of logical rules -- there are many other kinds, such as grammatical rules.

Most hierarchical systems I have seen have depended on some generative principle that applies to the relationships between all adjacent levels. If my analysis of perception is reasonably correct, evolution doesn't work that way. A new level is added only when the type of computing process typical of the highest existing level, elaborated horizontally and perhaps vertically too as far as possible, is no longer sufficient to preserve the organism. The next true level that is added is of a different type, using a different principle for computing perceptions.

I would have said all this earlier if I had thought of it.

Best, Bill P.

Date: Sun Aug 23, 1992 3:52 pm PST
Subject: Re: agre, facism

I'm testing the follow-up feature of the rn command. Hope this works.

In article <92Aug22.140417pdt.38019@huh.parc.xerox.com> sibun@parc.xerox.com writes:

>i really recommend preston's paper. it starts at exactly the same
>place as yr blind men paper, adn draws a v. diff conclusion.

I send a preprint request to Beth Preston but I am just burning with curiosity. Did she start off with the simultaneous equations for a closed loop system and come up with different conclusions about its apparent "behavior" in relationship to environmental variables and constraints?

By the way, I found that I do have a copy of a paper by Agre and Chapman (it's in the "Designing Autonomous Agents" collection edited by Pattie Maes, MIT, 1990). I have read it now. I see that there are some sentences that seem consistent with PCT. For example, on p. 21 they say "a plan does not directly determine an agent's actions" -- which is precisely true from PCT perspective; it directly determines an agent's PERCEPTIONS (assuming that a plan is a specification of intended outcomes, which seems to be exactly what A&C are trying to say).

They do make some, well, strange statements, the meaning of which seem to be left to the reader's imagination. For example, I never got a clear explanation (I thought) of this gem (also on p. 21): "We have found, in the case studies that we conducted, that participating in the flow of the environment, rather than attempting to control it, can simplify the machinery required to account for the organization of activity". If they are not saying something trivially obvious I would really like to know what it is. As it stands, it sounds to me like they are saying that the internal organization of, say, a stick that falls when you stand it on end and release it is simpler than that of a stick that maintains its vertical orientation when it is released. Is that it???

Best regards Rick

Date: Sun Aug 23, 1992 4:45 pm PST

[Avery Andrews 920824.1034]

(Bill Powers 920822.0800)

Thanks for the pointers. Two further observations.

First, I am very sceptical that people normally perceive the wheel angle. In fact, once I read an article in which you were advised to put some tape on the wheel at the position that is at the top when the wheels were pointing straight ahead, in order to make the wheel-angle information available. If people are being advised to set up visual aids to make it more available, that suggests that they aren't normally using it. So my guess is that rate of movement and net changes in angle (however observed), are what is actually noticed.

Second, while taking a long drive yesterday I noticed something that doesn't fit into my story as told, which is that on long distance drives at least, most of my (and my wife's, as far as I could see) steering movements were quick & small adjustments, whereby the wheel was turned various distances, but at what seemed subjectively to be a pretty uniform rate. These motions seemed to be completed before there was any noticeable change in the heading of the car. I don't yet have a real story about what's going on here, but I think it involves perceiving the car-heading to be wrong, ordering up more-or-less enough path-curvature to change it quickly enough, and then repeating this to straighten the car out again.

Avery.Andrews@anu.edu.au

Date: Sun Aug 23, 1992 6:51 pm PST
Subject: Re: agre, chapman

(ps 920823.1900)

In article <92Aug22.140417pdt.38019@huh.parc.xerox.com>
sibun@parc.xerox.com writes:

>i really recommend preston's paper. it starts at exactly the same
>place as yr blind men paper, adn draws a v. diff conclusion.

I send a preprint request to Beth Preston but I am just burning with curiosity. Did she start off with the simultaneous equations for a closed loop system and come up with different conclusions about its apparent "behavior" in relationship to environmental variables and constraints?

no; you saw the abstract. she started off w/ saying that behaviorism and cognitivism are looking at different aspects of the same thing.

They do make some, well, strange statements, the meaning of which seem to be left to the reader's imagination. For example, I never got a clear explanation (I thought) of this gem (also on p. 21): "We have found, in the case studies that we conducted, that participating in the flow of the environment, rather than attempting to control it, can simplify the machinery required to account for the organization of activity". If they are not saying something trivially obvious I would really like to know what it is. As it stands,

they are talking about ``leaning on the world.''

cheers. --penni

Date: Sun Aug 23, 1992 9:55 pm PST

Subject: From Intensities to Sensations

[from Gary Cziko 920824.0010]

Bill Powers 920822

>Consider sensations versus configurations. A sensation is a weighted
>sum of intensity signals. So the number of possible sensations is the
>number of different and independent weighted sums. If there are nine
>intensity signals, there can be nine independent sensation signals at
>a given time. Let's suppose that these nine sensations are
>discriminable tones of a scale. Clearly, only nine notes at a time can
>behave as independent tones.

Perhaps I need to go back to _BCP_ again, but I don't quite understand this. And I get especially nervous when I don't understand something which you preface with "clearly."

With nine intensity signals, why does "nine independent sensations signals at a given time" mean? Do you mean nine sensation signals at the same time? With nine intensity signals ranging from zero to some maximum neural frequency it seems to me that there would be a lot more than just nine sensations possible just using weighted sums. Consider at all the colors we perceive with only three (I think) intensity signals related to color. What am I missing here?

Once I understand this, I may have other questions about the perceptual hierarchy raised by your interesting post to Martin Taylor.--Gary

Date: Mon Aug 24, 1992 8:00 am PST

To Ed and anyone else who has comments on the topic,

At the conference you mentioned to me how "an unhappy past does not intrude on a happy present." This made sense, but at the same time I was hesitant to accept it given that it most therapists don't seem to think that way. Still not convinced, I think about the fact that many control systems are operating below conscious awareness. Many "autonomic functions" and probably many other systems are operating on other issues and problems (should I stay in school or not) even when we are attending to some other problem (should I call Marcy or Amy). Some of these systems are not attended to probably because the error is too great to be aware of. Now I don't understand why attended error is worse than unattended error, but don't you think that "working through the past" works (sometimes) because it somehow reduces or eliminates the error that is always there. I can kinda start to see how various unresolved psychological stresses would result in specific health problems (which the doctors never figure out), for the "body" is experiencing stress even though the "person" doesn't know it.

Certainly, I agree that one should not attempt to fix what one cannot control, but doesn't some of the psychoanalytic therapy attempt to take the child's perceptions of an event and make them into adult perceptions? It seems to me that PCT explains why working through the past is necessary, and how it should proceed (change perceptions so that error no longer persists).

Mark

Educational Psychology 210 USmail: 405 South 6th St. #4
College of Education Champaign, IL 61820
Univ of Illinois at Urbana-Champaign
phone: (home) 351-8257 e-mail: (Internet) m-olson@uiuc.edu
(office) 244-8080 (Bitnet) FREE0850@uiucvmd

Date: Mon Aug 24, 1992 8:01 am PST
Subject: Re: Hierarchies, degree of freedom

[Martin Taylor 920824 12:00]
(Bill/Mary Powers 920822)

Since Bill copied part of his mail to me to the net, I'll copy part of my response, also slightly edited:
================

I can't agree that a change in combination method alters the NUMBER of available degrees of freedom. It only changes the shape of the (hyper)surface that defines variation in value for any one degree of freedom.

The number of degrees of freedom available is computable at any instant, but things become harder if you include time variation, because then the effective degrees of freedom depend on the tmeporal characteristics of the feedback loop. Even in a linear system that can be hard (the first time I came across fractional df was in Blackman & Tukey's Measurement of Power Spectra, 1958). How many df per second there are depends not only on the bandwidth, but also on the relative power in each delta frequency band. When the system is nonlinear, the equivalent spectrum (if you can talk about such a thing) may be very dependent on amplitude, and in a logical system it can depend on all sorts of (apparently) trivial contextual factors. So when you talk about the increase of df for sequence variables, you are talking about summing over the duration of the sequence the df inherent in all the control loops at any one level over that time period.

The one thing you cannot do is increase the available degrees of freedom. Every loop has a bottleneck somewhere. A control structure cannot control more df than the bottleneck permits. All I need for my argument is that the bottleneck is not in the sensory input structure, and may be in the muscular (and glandular, etc.) output systems. If it is in the outer world behaviour, that's fine. It just must not be in the perceptual input system.

Configurations, no matter how they are defined, cannot produce degrees of freedom. If there are nine notes from which configurations can be produced, there can be no more than nine independently variable ways to produce different configurations. Of course, if each note is discriminable as on or off, there are 2^^9 different configurations, but that's not to say there are 2^^9 degrees of freedom for configuration, because you can't have so many all independently variable. If you turn one note from on to off, you change the configuration you had, as well as the one you now have. They work together as aspects of the same single degree of freedom. How they work together is contingent on the states of the other 8 notes, but only on that. There are still no more than 9 df in the set of configurations.

As soon as you get to transitions, you get into the time-variation area I was talking about. For the 9 notes, and 2 time samples sufficiently separate, you get 18 df. No problem. But here the sensory systems overwhelm the muscular bottleneck even more strongly. Sensory bandwidths are on the order of tens to thousands of Hz, whereas muscular bandwidths are on the order of single-digit Hz or less. So my argument becomes much stronger.

>So I think that the degrees-of-freedom argument based on repeated application
>of weighted summations to previous linear combinations of weighted summations
>does not apply to the levels of perception as I have defined them. This is
>because with each new level, as I have defined levels, there is a change of
>logical type; no new level of perception is commensurate with any lower
level.

As you can see from the above, the df argument does not come from weighted sums. I quite accept the change of type at each level, though there may be some argument about whether a change of type means a change of mechanism. But that's for later.

Martin

Date: Mon Aug 24, 1992 8:38 am PST
Subject: From Intensities to Sensations

[From Rick Marken (920824.0930)]

Gary Cziko (920824.0010) asks
> With nine intensity signals, why does "nine independent sensations signals
at >a given time" mean? Do you mean nine sensation signals at the same time?

Yes; Nine linear combinations of nine variables can be INDEPENDENTLY and simultaneously brought to nine different values. It's just basic linear algebra -- you can solve, at most, N simultaneous linear equations when you N variables (and, of course, N unknowns).

> With nine intensity signals ranging from zero to some maximum neural
frequency >it seems to me that there would be a lot more than just nine
sensations possible >just using weighted sums.

Yes - there are an infinite number of ways of combining the nine inputs linearly to produce sensation signals. But you can only control nine of these linear combinations simultaneously (that's why there are only six sensations in my excel spreadsheet -- because there only six intensities; but you can still find sets of 6 linear combinations that are not orthogonal (they are linearly dependent) so you don't really have 6 independently controllable sensations. This is how you can have conflict at level N+1 even though there are plenty of available inputs at level N).

penni sibun says (re:Preston paper):
> you saw the abstract. she started off w/ saying that behaviorism and
>cognitivism are looking at different aspects of the same thing.

Well, memory is the first to go (luckily). When did I see the Preston abstract? Is it in an earlier post. It does sound like she is saying what I am saying (behaviorism and cognitivism are looking at different aspects of the same thing -- but what is that thing, according to Preston?)

Best regards Rick

Date: Mon Aug 24, 1992 11:12 am PST
Subject: Degrees of freedom; Unhappy pasts

[From Bill Powers (920824.1100)]

Avery Andrews (920823) --
> First, I am very sceptical that people normally perceive the wheel angle.

In HPCT, all that is required in order to say that a perception exists is that a neural signal exists in a perceptual pathway. This has nothing to do with consciousness. In a spinal control system, the signals standing for muscle stretch and tendon tension are present at all times, at some magnitude. Copies of those signals rise to the brainstem, the cerebellum, the midbrain, and (by some direct pathways) to the motor cortex. The control systems involving those perceptions are always active. But we are seldom conscious of those signals unless we deliberately attend to them, or something goes wrong that draws our attention to those levels of organization.

Second, while taking a long drive yesterday I noticed something that doesn't fit into my story as told, which is that on long distance drives at least, most of my (and my wife's, as far as I could see) steering movements were quick & small adjustments, whereby the wheel was turned various distances, but at what seemed subjectively to be a pretty uniform rate. These motions seemed to be completed before there

>was any noticeable change in the heading of the car. I don't yet have a real
>story about what's going on here, but I think it involves perceiving the
>car-heading to be wrong, ordering up more-or- less enough path-curvature to
>change it quickly enough, and then repeating this to straighten the car out
>again.

There are other controlled variables at lower levels, important ones being the senses of sideward and rotational acceleration that indicate the start of a movement of the car. If something accelerates the car to the left (a bump), you feel an acceleration to the left as the car presses sideways against you, accelerating your body toward the left. You immediately turn the wheel to the right, reducing your body's sideward acceleration (which you feel as a force). This doesn't completely prevent the car's direction from changing, but the higher- level systems based on vision can correct the residual error. So one of the lower-level reference signals that is set by the higher-level driving system is "zero lateral acceleration."

An interesting demonstration of this sort of effect shows up when you accelerate the car with the foot-pedal. At the same time that the speed-control system increases the reference level for visually- detected speed, it raises the reference signal for forward acceleration, which you feel as your stomach muscles tighten to force your trunk forward and thus keep it stationary relative to the seat and steering wheel. If the transmission happens accidentally to be in neutral, the reference signal for greater forward force causes your body to pitch forward, as if someone had slammed on the brakes.

Normally we are conscious as though from the viewpoint of a higher- level system -- which level depends a lot on the level you habitually adopt as a point of view. People like us spend a lot of time doing logic and verbal manipulation, so very often we are unaware of the workings of the lower level control systems. They operate, however, just as well without awareness, and possibly better.

So you're right in being skeptical about people perceiving wheel angle. If, however, you asked them what the wheel angle was, they could tell you by paying attention to the positions of their hands or by attending to the part of the visual field where the steering wheel is. This doesn't mean that wheel angle isn't being controlled even when they're not attending. If you're the passenger, just reach out and tug at the wheel -- you'll feel resistance from the driver even before the car has begun to deviate, and before the driver yells at you. If you keep your disturbance small, in the range of normal disturbances that arise from little irregularities in the road, the driver might not even notice -- but the disturbances will still be resisted. That tells you that wheel angle or at least angular velocity is under active control by a system that's currently not in awareness, but is still a necessary level in the steering hierarchy.

I'm delighted that you're thinking about control theory while you drive. You'll learn a lot more about it by sorting out real control experiences than you will writing model programs. Just don't forget to attend to the higher levels every second or so!

You'll notice an implication that we can attend to perceptual signals that are not at the highest levels in the brain. I'm not sure that's true, but it seems to be true.

Gary Cziko (920824.0010) --

This one seems to have kept you up pretty late.

> I get especially nervous when I don't understand something which you preface
>with "clearly."

Roger. I get the same feeling in reading philosophy when I run across italicized words. Invariably, these are the words most in need of definition and most lacking in it. Example: intention is a sense of _aboutness_.

With nine intensity signals, why does "nine independent sensations signals at a given time" mean? Do you mean nine sensation signals at the same time? With nine intensity signals ranging from zero to some maximum neural frequency it seems to me that there would be a lot more

>than just nine sensations possible just using weighted sums.

Yes, there could be jillions of different sensation signals. But only nine AT A TIME can be controlled with respect to independently-set reference levels without creating conflict. This is because independent control amounts to solving simultaneous equations. You have nine (perceptual) functions of the nine intensity signals that must have specified values at the same moment. With nine equations in nine unknowns, a solution is possible if the perceptual functions are linearly independent. If you try to control 10 or more functions of 9 variables, you won't find a solution.

More than nine sensation-signals can be present at the same time, even with only nine different intensity signals at the level below. But only nine of them at a time can be controlled relative to specific referenc levels.

Consider at all the colors we perceive with only three (I think) intensity signals related to color. What am I missing here?

This has worried me, too. If we just look around at the environment, we see an incredible number of hues of color. Is there a control system for every hue? Part of the answer is in the fact that when we ADJUST color (for example, with a TV set's color controls), we attend to only one place in the visual field, around the center of vision. It's as if we can exert active control only for what is in the center of attention. All the other perceptual variables at the same level just sit there in whatever state we left them (if they'll stay that way -- otherwise they drift). When a lot of variables of the same type need controlling simultaneously, you get the one-armed paper hanger effect (OAPHE). Our attention jumps around among the variables, and we have to switch our outputs from one reference signal to another, trying to keep them all appropriately set as multiple disturbances upset all the variables.

This is clearly an important part of an HPCT model that doesn't exist in the present form. In a way it makes the modeling job easier -- it says that we don't need a separate control system for each parallel instance of a given type of perception. We seem to be able to keep some small number -- 7 plus or mine 2? of control processes running in parallel at the higher levels, at least within consciousness, but not more. At the lower levels, we can apparently run far larger numbers at the same time, but mostly because we don't need to be aware of them.

At or above some level, it seems that we can set up control processes to run at the same time, but without attention they tend to decay and drift. It seems to me that this is an area where we need a lot of experimental data. We can measure control parameters rather easily and quickly. Wouldn't some research group like to investigate how control varies with attention? There are gobs and gobs of vital information to be obtained here. This is a high-priority subject. Before we try to construct models that can do this sort of "scanning" process, we have to know what phenomena need modeling.

Mark Olsen (920824) --

RE: unhappy past.

The past doesn't have any effect on the present, any more than the future does. The past is gone, the future isn't happening yet. If an unhappy childhood has any effect on an adult person, it's because of attitudes, beliefs, opinions, and ways of doing things that exist RIGHT NOW.

Memories, of course, are part of present time. Whether the content of a memory is a true recording of past experiences is irrelevant; what remains now is what constitutes the pool from which reference signals can be drawn. Ed Ford says to replace the unhappy memories with happy ones. Enlarge the pool and add to it memories that will be useful now, in the present world, in relation to the people with whom you now interact.

As you say, psychoanalysis does dwell extensively on the past, at least in theory. but psychoanalysis is an incredibly inefficient mode of therapy, and in many cases I know about the chief accomplishment is to teach a client to describe his or her problems in a specialized language, while becoming reconciled to having them go on forever.

Martin Taylor (920824) --

Yes, it seems best to go public on this debate, as it pertains to what's happening on the net.

I think that when you shift types as you go up levels, the degrees-of- freedom problem takes on some new aspects. You already showed how you could get 18 degrees of freedom out of 9, just by adding first derivatives. Couldn't we get 9 more by adding second derivatives, and so on? The question now is, what ELSE can be added that takes us out of the domain defined by our original concept of intensities -> sensations -> configurations?

Bandwidth really isn't the problem here -- we could keep all variations well within the available bandwidth and still have a problem to solve.

Consider this. We start with nine functions of nine variables, the value of each function being adjustable to any reference level, and each reference level now being adjustable with any time rate of change, independently (well within the maximum bandwidth, remember). So we could have r1 varying as sin(t), r2 varying as sin(2.7*t), and so on. All nine reference levels can thus vary independently with respect to oscillation frequency.

In addition, the amplitude of oscillation can be varied independently for all nine systems. There's an outer envelope set by bandwidth; the bandwidth and momentary frequency sets the maximum rate of variation in amplitude, but within that envelope, there is complete freedom. Now we have nine more dimensions.

At the event level, we can now have System 1 execute one complete oscillation with a given amplitude, System 2 execute one, and so on, to construct an event like a wave traveling one time across the systems. Or we could divide the nine reference signals into three groups and have a wave pass through groups 1,2, and 3 independently, choosing an arbitrary pattern for each one. So the same nine systems can be used to construct an infinity of different patterns of variation at the event level. We now have not only nine derivatives (at least) at an instant, but a time-spanning characteristic called amplitude, and a time-spanning pattern characteristic that can extend indefinitely through time.

I think it gets very difficult to define what we mean by degrees of freedom when we get beyond the first four levels. In one sense, we could say that the event level introduces just one new degree of freedom, because we can execute only one pattern at a time using the lower nine-component system. But suppose that at the event level, the complex event is a one-bar phrase of eight notes -- a little fragment of melody. At any time, only one phrase can be generated. But we can then generate a number of other phrases holding the first note constant, another set holding the second constant, and so on. Since we're talking about time-spanning perceptions now, we can compare a phrase emitted at one time with a phrase emitted at another time. We can differentiate one song made of one set of phrases from another song made with a different set of phrases -- even though they don't coexist. We've now created a conceptual space in which various components of a perception can change while holding other components the same. We can define axes arbitrarily in such spaces, can't we?

Behind the strict degrees-of-freedom problem as we originally conceived it, I think there are some level-specific assumptions that aren't obvious. One is the concept of simultaneity -- we say that the nine systems must in effect solve _simultaneous_ equations. When we think of levels of perception in which temporal patterns are the variable, the idea of simultaneity no longer applies. We can solve a control problem that requires the nine systems to achieve their reference states ONE AT A TIME or IN A CERTAIN PATTERN or IN CERTAIN SPACE-TIME RELATIONSHIPS and so on. These are actually solutions to conflicts that would arise if we demanded that all nine systems reach zero error at once. These solutions may explain why higher levels exist.

Even if we think of only nine letters on a keyboard with one finger for each letter, typing the word "keyboards" is impossible to do with a linear combination of finger-presses. We can clearly control each finger simultaneously and independently with nine control systems, but when we try to do so, the keyboard won't respond and we won't see the word "keyboards" on the screen. We'll see whichever letter was hit ahead of the others by a millisecond. The only way to solve this problem is by introducing a time-spanning event in which the reference signals for each finger are specified one at a time. A good typist can learn to do this at the event level; a beginner does it at the sequence level, and much more slowly.

And with only the nine fingers plus two more degrees of freedom (x and y), we can type everything that it is possible to type. How many degrees of freedom have I used in saying all this so far?

Date: Mon Aug 24, 1992 12:10 pm PST
Subject: astro in C

The C version of astro is beginning to exist (the control circuits are there, but not the reorganization), but unless people express an interest in playing with it I will probably let it languish indefinitly.

Avery.Andrews@anu.edu.au

Date: Mon Aug 24, 1992 12:11 pm PST
Subject: Re: Degrees of freedom

[Martin Taylor 920824 15:40] (Bill Powers 920824.1100)

I think that when you shift types as you go up levels, the degrees-of-freedom problem takes on some new aspects. You already showed how you could get 18 degrees of freedom out of 9, just by adding first derivatives. Couldn't we get 9 more by adding second derivatives, and so on? The question now is, what ELSE can be added that takes us out of the domain defined by our original concept of intensities -> sensations -> configurations?

The answer to the last question there is "Nothing." At least in respect of the number of available degrees of freedom. Certainly adding a second derivative gives you 9 more, but you need another independent time sample to get it. You can get as many degrees of freedom as you want by taking sufficiently many independent time samples. Bandwidth really isn't the problem here -- we could keep all variations well within the available bandwidth and still have a problem to solve.

Well, in a sense it is THE problem, but when you get non-linear you can't interchange time and frequency freely, so you need some kind of surrogate for bandwidth. But for heuristic discussion, bandwidth will do. All nine reference levels can thus vary independently with respect to oscillation frequency. In addition, the amplitude of oscillation can be varied independently for all nine systems. There's an outer envelope set by bandwidth; the bandwidth and momentary frequency sets the maximum rate of variation in amplitude, but within that envelope, there is complete freedom. Now we have nine more dimensions.

No you don't, if I understand your layout. You only have the selection of values within the original nine dimensions. You do change the information rate if you change the power level (SNR) within a dimension, but you don't add any degrees of freedom.

So the same nine systems can be used to construct an infinity of different patterns of variation at the event level. We now have not only nine derivatives (at least) at an instant, but a time-spanning characteristic called amplitude, and a time-spanning pattern characteristic that can extend indefinitely through time.

Sure. There are an infinity of possible settings of any single variable, too. But that variable would have only one degree of freedom instantaneously, and 2*W degrees of freedom per second, where W is the equivalent rectangular bandwidth of the variation of that variable (in a linear system).

[Discussion of making melodic phrases, which can be compared with ones in memory] We've now created a conceptual space in which various components of a perception can change while holding other components the same. We can define axes arbitrarily in such spaces, can't we?

Sure, but no more informationally independent ones than the original set of notes allows- 2NWT where N is the number of simultaneously hearable notes and T is the duration of the phrase. (I assume chords are allowed in your phrases). You can define an infinite number of different axes, but measurements on one won't be independent of measurements on another.

(Parenthetically: I think exactly this does happen in any "real" control hierarchy, and it is what allows the neural architecture to deviate from strictly reciprocal connections in which reference links parallel perceptual links between higher and lower ECSs.)

When we think of levels of perception in which temporal patterns are the variable, the idea of simultaneity no longer applies. We can solve a control problem that requires the nine systems to achieve their reference states ONE AT A TIME or IN A CERTAIN PATTERN or IN CERTAIN SPACE-TIME RELATIONSHIPS and so on. These are actually solutions to conflicts that would arise if we demanded that all nine systems reach zero error at once. These solutions may explain why higher levels exist.

Yep, that solution to the problem is directly implicit in the degrees of freedom argument. If there were no ouput bottleneck in degrees of muscular freedom, the patterns WOULD be simultaneously achievable and there would be no conflict. It is the limitation of available df for control that leads to all conflict situations (or most, anyway; I'm not prepared to stake my reputation on "all").

Even if we think of only nine letters on a keyboard with one finger for each letter, typing the word "keyboards" is impossible to do with a linear combination of finger-presses. We can clearly control each finger simultaneously and independently with nine control systems, but when we try to do so, the keyboard won't respond and we won't see the word "keyboards" on the screen. We'll see whichever letter was hit ahead of the others by a millisecond.

So here, the df bottleneck is in the external world. The keyboard has one instantaneous df, and as many df per second as it will accept keystrokes. And with only the nine fingers plus two more degrees of freedom (x and y), we can type everything that it is possible to type. How many degrees of freedom have I used in saying all this so far?

I find it easier with ten. But I don't know how many df you have used. A rough estimate would be to count the letters, but that could be either an understimate or an overstimate. It could be an underestimate because you may have done a lot of erasing and retyping, which does not show in the post, and it could be an overestimate because of the sequential redundancies among the letters, which remove from you the choice of which key to strike on some occasions if you want to make sense. But the number of characters is a reasonable order-of-magnitude estimate.

Martin

Date: Mon Aug 24, 1992 12:42 pm PST
Subject: Re: relation of effort to result

[Martin Taylor 920824 16:30]
(Bill Powers 920821.0800) (catching up a little)

The only way it can seem that gross effort corresponds to gross result is by confusing the effort with the result. We simply assume that if an effect is going in a given direction, the efforts we feel must be aimed in that direction, too. This is what I mean when I say we name behaviors after their controlled results, not after the efforts required to keep them controlled.

When Astro is making a final approach to Mother in more thabn one dimension, its jets are firing opposite to the direction of travel. When its path is curving and its velocity is changing, its jets could firing at any angle to the direction of travel.

I think I may have already mentioned an article in Science that relates to this: "The motor cortes and the encoding of force" by Georgopoulos, Ashe, Smyrnis, and Taira, Science June 19, 1992 p 1692. They showed that the motor neurons in the monkey cortex encode not the force being applied by the muscles, but the force needed to counter the (visual) disturbance added to the location of a target on which a cursor was maintained in the presence of a constant bias force that was already being compensated. The total force direction could be anything, depending on the bias, but the encoded force was toward the target, countering the disturbance.

Martin

Date: Mon Aug 24, 1992 1:00 pm PST
Subject: great new taste sensations

(ps 920824.1400)

> [From Rick Marken (920824.0930)]
>
> abstract? Is it in an earlier post. It does sound like she is saying what I am saying (behaviorism and cognitivism are looking at different aspects of the same thing -- but what is that thing, according to Preston?)

she doesn't give it a name. she characterizes some of the commonalities and contrasts interactionism with those commonalities.

--penni

Date: Tue Aug 25, 1992 4:19 am PST
Subject: Re: steering

[From Chris Malcolm]

There is an interesting control observation to be made concerning motorcycle steering. A motorcycle can only stably turn a corner while leaned over so as to counteract the centrifugal force. The geometry of motorcycle steering is so set up that when the bike is travelling forwards and leaning over it "wants" to turn its wheel appropriately for the turn; though some of them are better than others at this. The problem is how to start turning, since if you simply turn the bars in the appropriate direction the bike will fall over to the other side. The fastest method is simply to turn the bars in the opposite direction to the required turn. This causes the bike to fall over rapidly in the required direction. When the bike has fallen over enough, the fall is "caught" by turning the bars in the required direction of travel, thus producing the balancing centrifugal force. In practice the bars usually do not require to be turned by force, merely permitted to line up as they wish to, i.e., the turning of the bars into the turn is accomplished by relaxing the force applied to produce the initial counter-steer.

This initial turning of the bars the "wrong" way is known as "counter-steering". What is interesting about it is that most experienced motorcyclists have learnt to do this, but that, unless they have been exposed to explicit training, not only do they not realise that this is what they do, but they usually strongly assert that they do the opposite. Frequently considerable practical experiment is required to disabuse them of their wrong hypothesis, and permit them to correctly observe what they really do. Some find it impossible to learn this perception, and continue to assert strongly that they do not do what a simple video can show them to be doing.

Now if a motorcyclist actually starts a turn by counter-steering, but strongly asserts that they do exactly the opposite, needing quite a lot of practical experiment to manage to observe correctly what they do, this means not only that they learnt this skill without being conscious of what they were doing, but were able to learn and practise one behaviour _while_ believing that they were doing the opposite. In fact, having had this experience myself, I can attest that oddly enough one has to _learn_ to perceive how one is turning the bars, and that this is difficult, despite the fact that when rapidly changing direction one exerts considerable force to apply the counter-steer.

One of the difficulties in doing this learning is that it is very difficult indeed to experimentally impose variation on the arm movements demanded by these learnt skills. For example, try riding a bike and giving a capricious twist to the bars. It requires a supreme effort of will, beyond many people's capability. And what you cannot will it is difficult to become conscious of.

Clearly these apparent paradoxes arise because the "folk" theories of how control is learnt and executed are wrong. What is required is a theory of control which can be applied to the phenomena of counter-steering in which the paradoxes become predicted. Can CSG supply this? Can the "conventional" theories of control supply this?

Chris Malcolm

Date: Tue Aug 25, 1992 7:59 am PST
Subject: D of F; monkeys; motorcycles

From Bill Powers (920825.0800)]

Avery Andrews (920824) --

I'd like to see if I can get your C code to compile and run on my machine. Could you send it to me?

Martin Taylor (920824.1540) --

>You can get as many degrees of freedom as you want by taking
>sufficiently many independent time samples.

I think we agree. The key to what I'm trying to say is probably in that fact.

In a neural pathway, where redundant signals travel in parallel, there's no coherence between various partitions of the signal, so it must be difficult to compute the equivalent number of time samples per second. But in the middle of the frequency range of a neural signal, the number must be rather large. For example, in perceiving a shape like a filled square, the number of parallel signals contributing to the sense of squareness must be quite large (depending on the size of the square). If the square is shrinking or expanding, the derivatives could be computed over a pretty short time, I should think. Also, as Rick showed experimentally, the best definition for the "size" of a square is its area, so shrinking and expanding size at a uniform rate would mean increasing and decreasing the number of contributing channels at the square of the rate. What I'm getting at is that what we casually label "a" signal contains a lot of detail that isn't necessarily evident when we speak of something like "a sensation signal." These details can change greatly without altering the meaning of the signal at the sensation level.

>>Now we have nine more dimensions. [adding frequency and amplitude for
>>each of the nine reference signals]

>No you don't, if I understand your layout. You only have the selection
>of values within the original nine dimensions.

I'm not really arguing here, just trying to get my thoughts straight on this subject. Let me try to illustrate the point I'm stuck on without getting formal about levels. I'll go along with your ten variables instead of my nine, which number was picked because I had in mind an example with three groups of three.

Suppose the 10 variables (which we agree are controlled at the sensation level) are 10 notes on a piano, controlled by striking the keys with 10 fingers. There are certainly 10 degrees of freedom for controlling the loudness of these notes. Plotting a point in 10- dimensional loudness space specifies a loudness for all of the 10 notes. So this should use up all the degrees of freedom available.

Now we can add one more dimension: repetition rate of each note. And another: rate of increase of repetition rate. And another: rate of increase of loudness as the repetition proceeds. (This does not add 10 more dimensions for each, as I said in my last post -- only one for each). It's not likely that any real person could exert independent control of all ten variables, let alone of all 13 (!) dimensions, but we're talking "in principle" here.

Now where did these three extra dimensions come from? We both know that you can't control 13 functions of 10 variables independently, yet here we are doing that. The answer is that we didn't start with 10 variables: we started with 13. It's just that the sensation level of perception doesn't register three of the aspects of the lower-level signals _as independent aspects of the signals_. For a sensation- control system, all that matters is the instantaneous amplitude of the signal averaged over all redundant pathways. What makes the difference between one sensation signal and another is the weighted-sum function that creates different signals using weighting schemes that are not linearly dependent. There is no scheme at this level that differentiates signals on the basis of their time characteristics, patterns of change, and so on.

Nevertheless, a system that DOES differentiate signals on these bases will see that such variations are present in the original information -- that there are, in fact, more degrees of freedom in the sensation signals than the sensation-level systems can discriminate or control. You handle this by speaking of time-sampling. A sensation-level system, however, can't appreciate or control the variations in its own perceptual signals that a time-sampling would reveal. In your way of handling time, you speak from the objective viewpoint. For the system itself to make any use of the temporal information, however, it must contain perceptual functions that are sensitive to temporal aspects of the sensation signals. These aspects are passed through the sensation level without making any difference to the sensation-level control systems. And the sensation reference signals can be varied in ways that have no meaning to a sensation-controlling system.

Somewhere around 1956 or 1957, Kirk Sattley (hello, Silent One) suggested to me that what I was calling levels of perception might actually be dimensions of perception, not arranged in strict levels but constituting a space in which perceptions are controlled. I have opted to stick with the levels concept, in large part because of the physical arrangement of the nervous system and considerable information about identifications of types of perceptions with locations lower or higher in the brain. But I haven't forgot Kirk's proposal. In some cases it seems a better one than mine, especially those where we seem to use a high-level control process to achieve a low-level end (although that's often just a manner of speaking). I'm not ready to adopt it, however, because of the clear dependence of some kinds of perceptions on the existence of others, which seems to imply the necessity of hierarchical control.

It's profitable to think of the levels as being dimensions of the perceived world, because that's how we experience it. We don't experience the world in layers; all the levels exist in it at the same time, all mixed together, all in the same space. What I'm getting at is that these levels may (to make an epistemological guess) reflect more and more subtle aspects of what Wayne Hershberger called the "immanent order." If that's the case, then in 10 sensation signals we have actually 10 little packages of representations of the world, each package capable of far more subtle variations than can be appreciated by a perceptual function that responds only to weighted sums and acts only to control that one aspect of the world.

As I said, perhaps your idea of temporal sampling is the key to the added dimensions of experience that show up in the higher levels. Practically all of the higher-order perceptions are time-spanning, or time-irrelevant. At the sequence level, as long as the same sequence is progressing as it should (m-i-s-s-i-s-s-i-i-p-i) we experience a single thing, that particular sequence, with no error except where there's a mistake. As you prove that (x + y)**2 equals x**2 + 2xy + y**2 (Fortran), you're engaged in a single recognizeable rule-driven or logical procedure that remains the same until the job's done.

Temporal sampling at the lower levels (fourth) is probably done by hardware -- rapid-adapting neurons, local integral negative feedback, and so on. At higher levels, however, memory gets into the act, so that the time dimension can be exploited over very long spans of time. This allows introduction of a dimension such as "familiarity." Something is "familiar" if there's something like it already stored, regardless of how long ago it was recorded. You can say of a melody played on ten notes that it's like another melody, or that it's part of a longer piece, or that it was written in 1846, or that it's trite, and so on. A given melody can, independently, be all these things and many other things besides, all at the same time. So as you say, the dimensionality becomes unlimited -- even though everything has to filter through the ten note-perceivers.

As to the "bottleneck" effect, it certainly exists, but we have ways of getting around it to some extent. Pianos actually have 88 keys, while nearly all pianists have 10 fingers. Yet as my friend Sam says, there are always enough fingers to play the piece. This is reminiscent of my comments yesterday about the OAPHE (one-armed paper hanger effect). Because of the time dimension, we can multiplex (your term in the Paris paper) our outputs and make them look like many more than actually exist. A juggler can keep five balls in the air with two hands, because physical processes don't begin and end instantaneously. I'm beginning strongly to suspect that something similar happens in the brain: a control system of limited capability switches back and forth between multiple sets of lower-level systems, tweaking their reference signals (which persist for a while between tweaks) and going on to the next set, so a single system can behave like multiple systems controlling the same variable at the same time, only in different contexts at lower levels. A relatively small set of higher- level systems, therefore, can act like a much larger set. Again I plead for experimentation.

RE: monkey cortex (920824.1630) --

I saw the article, but I've given up on saving articles having something to do with control theory. There are too many. And it's just too much work to sort out the wild interpretations put on the data, as in

>" ... motor neurons in the monkey cortex encode not the force being
>applied by the muscles, but the force needed to counter the (visual)
>disturbance added to the location of a target on which a cursor was
>maintained in the presence of a constant bias force that was already
>being compensated. "

Motor neurons don't encode forces needed to do something else. They just produce forces by using muscles. This quote is an example of the tunnel vision that pervades medical-style research: label everything by its effects, and attribute the cause arbitrarily in the middle of the loop. To explain to Georgeopolis why the above statement is nonsense would require re-educating him from High School onward. Of course he's looking at part of a control system. But how could we ever explain that?

Chris Malcom (920824) --

I've also made this observation about bicycle riding -- the first thing you do when you turn left has to be to turn the front wheel right, to get the bicycle out from under you and initiate a bank.

Your main point is important. The hierarchy does not learn to control perceptions by describing them and reasoning about them. Describing and reasoning are the RESULT of learning to control, not the cause. A lot of people try to understand control theory strictly in terms of verbal descriptions of how they work -- and as you know, that's not good enough. The descriptions can be pertinent only after the right perceptions have been acquired. The way people describe their actions often has nothing to do with what's actually going on. Just think of people saying "You make me mad!" This is the main reason why people can't just read a book about PCT and understand it. They have to apply the words to real experiences; only then do they realize what the words actually mean.

-- -

Best to all, Bill P.

Date: Tue Aug 25, 1992 8:33 am PST
Subject: outa here

From Tom Bourbon [920825 -- pre-Alexander time]

Since I posted my message about the loss of funds at the lab, a number of you have responded with offers of (moral) support and a few have offered direct suggestions about employment, research funds and the like. Thanks, to all of you.

Yesterday, we learned that our salaries will continue for another three months. A bit of room to breathe.

I had wanted to jump into some of the discussion about "leaning on the world" -- is that a round about way of saying that tha actions of an organism combine with independent influences in the environment, so that the organism need only, indeed can only, add a small influence to what is going on? In that case, the nature of the controller becomes critical. Do the authors of work about "leaning on the world" address the nature of how control is achieved under those conditions?

I would say more, but hurricane Alexander is on the way to our island paradise, or so it seems as of the latest advisory. I am, as they say, outa here!

Tom Bourbon

Date: Tue Aug 25, 1992 9:07 am PST
Subject: CONTROL OF PERCEPTION!!!!!

[From Rick Marken (920825.0930)]

PCT has been around for quite some time. The basic idea of PCT is rather easy to state:

Organisms are closed-loop negative feedback systems that control various aspects of their own perceptual experience. Observable actions are the visible means by which organisms influence their own perceptions and compensate for those influences of the environment (disturbances) that would tend to move perceptions from reference levels specified by the organism itself. Thinking, planning, memory and imagination are pocesses that involve controlling perceptions without going through the external environment. Learning is the processes of reorganizing the existing means of controlling perceptions.

Basic PCT is still best summarized by the title of Powers' classic: Behavior: The control of perception.

We have been hearing a lot in the last year about many of the "hot" new approaches to understanding the behavior of living systems. Some are schools of thought, others are the approaches of individuals or labs. Examples are: Beer's bug, Artificial life, interactionism, Agre/Chapman Pengi system, Brooks' subsumption architecture (and general approach to robotics), fuzzy logic, chaos, dynamical systems, Artificial Intelligence, etc, etc.

What ALL these approaches (and many, many others) have in common is that NONE of them explicitly recognizes the basic organizing principle underlying the behavior of living systems -- that they CONTROL PERCEPTUAL INPUT VARIABLES. Thus, the essential insight of PCT is missing from EVERY ONE of these hot approaches. In fact, many times it seems like these investigators are going OUT OF THEIR WAY to avoid concluding that organisms control perceptions. This hit me while reading the Agre/Chapman paper. A&C said at lot of things that were consistent with PCT. For example, they correctly pointed out the problems for programmed output type models -- 1) unpredictable changes in the environment, 2) complexity of computing all required outputs [we'd call that the "inverse kinematics" problem] and 3) the unpredictable effects of planned activities (I'm being generous -- it sounds to me like A&Cs problems 1 and 3 are the same). So A&C state (albeit quite unclearly) the basic reason why programmed output systems will not produce life-like behavior -- [the reason. of course, is disturbances, in PCT terms]. They even say that behavior cannot be generated by a plan that "determines an agent's actions" -- which is EXACTLY what PCT says. But after all that they end up proposing a model that is described as one that generates actions -- when it is, in reality, one that generates perceptions. They just won't even SAY IT; the model acts to produce intended PERCEPTIONS.

A&C claim that their model doesn't control actions -- but they won't say what it controled What they do say is that the model has a mental plan that "serves as a resource that an agent can use in deciding what to do". How does the system use this resource? It turns out that it generates predicted outcomes (intended perceptions -- but they don't say it), it then generates actions and compares the actual outcome (as PERCEIVED) to the predicted outcome. The plan is a resource in the same sense that a reference signal is a resource -- it specifies an intended state of its own inputs.

So the A&C model (after all the verbal nonsense) is a control model -- and it acts to produce intended perceptual outcomes. But C&A don't realize this. The same was largely true of the Beer model and, as I recall, Beer seemed quite ademantly opposed to the suggestion that his bug was controlling perceptions.

Does it matter that these hot modellers don't understand that their models are controlling perceptions (when they are)?? Or is it just a verbal quirk -- PCTers like to say that behavior is the control of perception and hot research- ers like to say things like "behavior is the result of a communication process between plans and outcomes". I think the failure to even SAY that behavior is controlled perception IS a FUNDEMENTALLY IMPORTANT error -- not just a linguistic quirk. By failing to point to this important aspect of behavior (or by refusing to notice it it) those working on these hot approaches to behavior 1) fail to see all the great work that already has been done on modelling the control of perception 2) don't understand that the behavior of the system depends largely on how they design the perceptual functions 3) get lost with attempts to program output once the lower level perceptual control systems have been built and 4) don't understand how much of their own understanding of the situation is going into their modelling.

The moral: I just don't think there is much to be gained by trying to sift through the chaf of the hot approaches to understanding behavior in the hope of finding much wheat. Why do it anyway, when you have all the flour you need sitting right there in front of you with PCT?

The question: Why is there such reluctance on the part of those working on the hot approaches to behavior to even consider the possibility that behavior is the control of perception? Why do people in all these areas seem to ACTIVELY avoid even MENTIONING this as a possibilty? Has anyone ever seen the phrase "control of perception" used anywhere but in a paper by a PCTer -- whether it has been used correctly or not? What's the problem here?

Regards Rick

Date: Tue Aug 25, 1992 10:46 am PST
Subject: Re: D of F; monkeys

[Martin Taylor 920825 14:00] (Bill Powers 920825.0800)

We still have a conceptual problem about what is meant by "degrees of freedom."

>Suppose the 10 variables (which we agree are controlled at the
>sensation level) are 10 notes on a piano, controlled by striking the
>keys with 10 fingers. There are certainly 10 degrees of freedom for
>controlling the loudness of these notes. Plotting a point in 10-
>dimensional loudness space specifies a loudness for all of the 10
>notes. So this should use up all the degrees of freedom available.

at any instant. Yes.

>Now we can add one more dimension: repetition rate of each note. And
>another: rate of increase of repetition rate. And another: rate of
>increase of loudness as the repetition proceeds. (This does not add 10
>more dimensions for each, as I said in my last post -- only one for
>each).

You can't have a perception of repetition rate without having at least one repetition. For that, you have to have a duration over which the tone is on and a duration over which it is off, or an epoch at which a rapid rise gives way to a fall. You must be able to determine whether the the tone amplitude at time t+delta(t) is different from that at time t. If the range of possibly discriminable amplitudes at the second sample is the same as at the first sample regardless of what that value was at the first sample, then the second sample provides another degree of freedom. To get a repetition rate requires at least three independent samples, and probably many more (it would be three only at the fastest repetition rate that could be perceived as repetition rather than as a low tone). So to achieve the one degree-of-freedom description that is repetition rate, one needs several degrees of freedom for the amplitude of the tone, which is to say several independent samples of the tone. The same set of independent samples could be used to determine the rate of change of repetition rate, except that you need a minumum of 5 or 6. Each independent sample provides another degree of freedom.

So, if there are 10 notes, each with a possibly independent repetition rate, the number of degrees of freedom you need for amplitude evaluation is at a minimum 30, and probably many more, although each repetition rate number takes only one of those df, leaving quite a few for other parameters of the note amplitude envelope.

One probably does not perceive describable structures (trying to choose a neutral word) for most amplitude patterns. That, in my view, is a consequence of our inability to control as many df as we can sense, either instantaneously, or even more so, over time.
============

\>>" ... motor neurons in the monkey cortex encode not the force being
\>>applied by the muscles, but the force needed to counter the (visual)
\>>disturbance added to the location of a target on which a cursor was
\>>maintained in the presence of a constant bias force that was already
\>>being compensated. "
\>
\>Motor neurons don't encode forces needed to do something else. They
\>just produce forces by using muscles. This quote is an example of the
\>tunnel vision that pervades medical-style research: label everything
\>by its effects, and attribute the cause arbitrarily in the middle of
\>the loop. To explain to Georgeopolis why the above statement is
\>nonsense would require re-educating him from High School onward. Of
\>course he's looking at part of a control system. But how could we ever
\>explain that?

Sorry, it was my wording, not a quote from the original article. I am agnostic as to what the cortical neurons signal. They described them as being thought to encode force, and I guess I just took the words.

Why I introduced the article before and now again was because it seemed to me very clear evidence of the existence of a distributed control system demonstrated in actual neural output. The neurons respond with a signal that has nothing to do with the forces applied, but that relates directly to the error signal induced by the disturbance (the authors say that the neurons encode something they call "dynamic force" as opposed to the static bias force, but this sounds like a bit of waffle to me). The data seemed to me as direct a demonstration of the neural mechanism of control, as opposed to the perceptual fact of control, as one could wish. PCT studies normally demonstrate control by looking at the CEV that is presumed to provide the percept. Here we have not only that, but an apparent window into the control hierarchy itself. I don't see why that should bother you.

Martin

Date: Tue Aug 25, 1992 10:52 am PST
Subject: Submitting items to the Gopher server

If you have material which you want to submit to the CSG Gopher server, you can mail it to csg-news@biome.bio.ns.ca. It then goes into a folder called News in the CSG folder.

To use the Gopher server if you don't have a gopher client running, just telnet to gopher@biome.bio.ns.ca.

-- Bill Silvert at the Bedford Institute of Oceanography P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2 InterNet Address: bill@biome.bio.ns.ca

Date: Tue Aug 25, 1992 11:37 am PST
Subject: lotw

(ps 920825.1200)

i hope you get through the hurricane well!

> >From Tom Bourbon [920825 -- pre-Alexander time]
> I had wanted to jump into some of the discussion about "leaning on the world" -- is that a round about way of saying that tha actions of an organism combine with independent influences in the environment, so that the organism need only, indeed can only, add a small influence to what is going on? In that case, the nature of the controller becomes critical. Do the authors of work about "leaning on the world" address the nature of how control is achieved under those conditions?

i dunno, but i actually managed to find where agre introduces the term (``the dynamic strucutre of everyday life'', mit ai tr 1085, 1988, pp 26-27):

why does an understanding of an agent's interactions w/ its world lead to simpler hypothesies about its machinery? real agents _lean_on_the_world_ (emph in orig--ps). the world is its own best representation and its own best simulation. it isn't an obstacle or a problem, it's a helpful place. your interations w/ the world, both past and present, provide many ways to alleviate computational burdens. why conduct elaborate deductions about yr surrounding when you can look and see? in particular, why maintain elaborate control structures when you can look and see what needs to be done? why make highly detailed Plans when you can improvise? why required instant expertise when you can improve by just keeping on doing it? why try figuring it out yourself when you can collaborate w/ others who have been there? why insist on figureing out every situation afresh when you can trust yor accumulated experience? all the dynamic phenomena can work together to suggest imaginative ways to simplify machinery or even eliminate parts of it altogether. put someone in solitary connfinement and they fall apart: people rely on the _organizing_presence_of_the_world.

i am suggesting an inversion of values, not only in ai but in all forms of computational and psychological inquiry. faced w/ an empirical phenomenon to explain, our first explanatory recourse should be to dynamics, not to machinery. faced w/ a techinical problme to solve, our engineering should begin w/ dyunamics, not w/ machinery. heretofore, people have gotten prizes for inventing new machinery. but we've got far too much machinery. i would like to suggest that people get prizes for getting_rid_of_ machinery.

cheers. --penni

Date: Tue Aug 25, 1992 12:03 pm PST
Subject: Hotshots; encoding

[From Bill Powers (920825.1300)] Tom Bourbon, Rick Marken (920825) --

Rick: "The question: Why is there such reluctance on the part of those working on the hot approaches to behavior to even consider the possibility that behavior is the control of perception?"

Tom: " Do the authors of work about "leaning on the world" address the nature of how control is achieved under those conditions?"

It pains me, too. All of these people are fumbling around gradually discovering the facts of control, doggedly going down every blind alley and backing out again for another try. Chapman or Agre (I forget which) told me a month or two ago that Pengi is old stuff -- they've gone on to much bigger things now. So Pengi is a broken toy abandoned amid the mess of the playroom.

As I said to Rick on the phone a couple of days ago: Agre, Chapman, Brooks, and all the rest of that genre strike me as young hotshots cranked up to overload, working at top speed and hoping to get lucky and hit the bigtime payoff. They're brilliant, impatient, tremendously competitive, and undisciplined. It's probably hopeless to try to get them to settle down for a minute and learn something, like how control systems work. They'd rather figure it out themselves, because basically they don't think that anyone else could come up with anything as good as what they produce.

Right now all these people are going through the discovery phases that every CSGer has gone through, and that I went through starting in 1953. The difference is that they don't seem to want to learn control theory, so they're working with one hand tied behind their backs.

So what do we do about it? I don't think it's worth a lot of effort. They're going to resist all the way. If they do discover control theory on their own, they aren't going to come running to us with the news. It will be too embarrassing to admit that control theory existed before they were born, that all the engineers knew about it and could have told them if they'd just asked. To hell with 'em. There are plenty of people who want to know about control theory. I've wasted enough hours and days and weeks trying to make some sense out of the confused ramblings of dilettantes. Let's just get on with the work.

Tom, I hope there's something left of Galveston when you get back.

Martin Taylor (920825) --

Thanks for the added analysis of df. I bow to your superior experience with this field.

RE: the description of motor neurons encoding for the consequences of motor action. One of my objections was to the use of terms that don't mean anything, like "encoding." As far as I know, there's no special code in a motor output signal. It's just a signal that gets bigger and smaller. There's nothing about it that says what it's going to be used for. It just makes a muscle get tight.

To say that "... motor neurons in the monkey cortex encode not the force being applied by the muscles, but the force needed to counter the (visual) disturbance ..." makes it sound as if the motor neuron knows what's going to become of its output signal, or that it would give some special coding to the signal if it were used to counter visual distubances instead of just making a muscle get tense. In fact, a motor neuron signal creates a force, period. What the authors of the article discovered was that output does not correlate with outcome, but with disturbances. We already knew that. You can verify that without looking inside the brain. They did NOT find out how the brain manages to produce just the right signal to the right motor neuron to counteract the visual disturbance. They did not, in other words, find the control system -- only its output function.

>Here we have not only [the CEV being controlled], but an apparent window into
>the control hierarchy itself. I don't see why that should bother you.

That's not what bothers me. What bothers me is the attribution of a causal property to the motor neuron, an attribution that makes it seem as if the neuron could know that THIS time its output is being used to counteract a visual disturbance, instead of for example gravity. I'm not bothered by neural evidence that control systems really exist. I expect neural evidence to show that. I'm bothered by the implication that finding the motor neuron that produces the appropriate force is enough to explain the result. We know a force has to be produced if a disturbance is counteracted. We know that a neuron, or set of neurons, must produce it. What Geogeopolis did was to find the neuron. If he had gone on to raise the question of what generated the signal drove that neuron in such a singularly appropriate way, I might have saved the article.

Maybe I'm just getting burnt out.

Best, Bill P.

Date: Tue Aug 25, 1992 12:50 pm PST
Subject: control, trends

[From Rick Marken (920825.1330)]

First, a quick question to the net (especially Gary Cziko and Dennis Delprato):

Does anyone know of any written critique of Behaviorism which explicitly cites the following problem (that Gary mentioned some time ago) :if human behavior is controlled by the environment (as claimed by Behaviorists) then Behaviorists themselves should not be able to exert control over people (as they say that they can -- and should) because they are under control? What I want is a critique that points out that control cannot be exerted by agents that are controlled. I have never seen any detailed critique of Behaviorism from this perspective. The critiques I've seen usually just say "there is so a mind" to which the Behaviorists reply "is not" to which the cognitivists reply "is so" and so on -- a process called philosophical psychology.

A random note -- it seems that the trendiness in science is somewhat self-correcting. The following was posted to sci.cognitive:

> I am writing an article on AI Winter Refugees: un(der)employment
>among AI professionals.

Looks like ai has run its course. I guess that's why we're seeing new words -- artificial life might work for a few years (ai worked for about 7); neural nets should have peaked in a couple if they have the staying power of ai. I hope chaos and fuzzy logic don't create even larger unemployment lines. I just wish PCT could have it's shot at fame -- if it's just another trendy science we would at least have a few years of fun. But I think it would prove to be a wonderfully enduring surprise. Ah well.

Best regards

Rick

PS. Tom. Good luck with the hurricane. But I must admit that it's nice to see that at least ONE bad thing is not happening in LA.

Date: Tue Aug 25, 1992 1:10 pm PST
Subject: Modeling vs. WHAT?

[From Bill Powers (920825.1400)] Penni Sibun (920825) --

> why conduct elaborate deductions about yr surrounding when
>you can look and see? in particular, why maintain elaborate control
>structures when you can look and see what needs to be done? why make
>highly detailed Plans when you can improvise? why required instant
>expertise when you can improve by just keeping on doing it? why try
>figuring it out yourself when you can collaborate w/ others who have
>been there? why insist on figureing out every situation afresh when
>you can trust yor accumulated experience?

At one level I agree with all these rhetorical questions, as long as you don't apply them to PCT. At another, if you mean these questions generally, I think you're risking looking like a nihilist, or a naive realist.

These questions, at another level, imply accepting as an explanation of behavior what I would consider a statement of the problem. In fact, it doesn't look to me as though you're trying to _explain_ behavior or experience at all; you're just accepting it at face value with no explanation. You ask "why maintain elaborate control structures when you can look and see what needs to be done?" PCT is a model that attempts to explain how it is that you can look, see, and do. PCT explains how it is, for example, that two people can look at the same environment and see different things; how they can, even supposedly seeing the same thing, see that two quite contradictory things "need to be done," and how, even if they end up doing the same thing, they end up doing it by totally different means. PCT puts the "needing" into the organism, not into its environment. It interprets "seeing" as a process in the brain, perception, which entails as much construction as representation. It sees "doing" as a process of control. Are you really suggesting that we put all that intelligence into the non- living world? If so, I'd say that you have a mammoth job of explanation ahead, if you want anyone to believe you.

Or are you pulling someone's leg?

>faced w/ an empirical phenomenon to explain, our first explanatory
>recourse should be to dynamics, not to machinery.

I don't call that an explanatory recourse; it's just describing the dynamics of the phenomenon (usually from the standpoint of an unspoken theory). Once you've described that, the question of HOW IT WORKS arises -- i.e, the question of the machinery that accounts for these appearances. PCT is about how it works.

>i would like to suggest that people get prizes for getting_rid_of_ machinery.

Does this mean that it's modeling itself that you don't like? Do you think there's any point in talking about nerves, brains, and such machinery? I'm having a hard time figuring out where you're coming from.

Best, Bill P.

Date: Tue Aug 25, 1992 1:57 pm PST

[Martin Taylor 920825 17:30] (Bill Powers to Penni Sibun 920825.1400)

>> why conduct elaborate deductions about yr surrounding when
>>you can look and see? in particular, why maintain elaborate control
>>structures when you can look and see what needs to be done? why make
>>highly detailed Plans when you can improvise? why required instant
>>expertise when you can improve by just keeping on doing it? why try
>>figuring it out yourself when you can collaborate w/ others who have
>>been there? why insist on figureing out every situation afresh when
>>you can trust yor accumulated experience?
>
>At one level I agree with all these rhetorical questions, as long as
>you don't apply them to PCT.

I don't understand this answer. When I read Penni's posting, I wrote her privately that you yourself could not have described PCT better. If the "you" in her posting is an ECS or a group of ECSs, doesn't this descibe pretty succinctly what is going on in the control hierarchy? That's how I interpreted her posting. Either I understand PCT less than I thought I did, or one of us is grossly misunderstanding Penni.

>>faced w/ an empirical phenomenon to explain, our first explanatory
>>recourse should be to dynamics, not to machinery.
>
>I don't call that an explanatory recourse; it's just describing the dynamics
>of the phenomenon (usually from the standpoint of an unspoken theory).

But it is precisely the dynamics of the control hierarchy that describe the modelled behaviour. Again, I miss the point of your criticism, since it seems to me that you and she both have the same intention of simplifying mechanism to a point at which it can be directly tested through the dynamics.

Penni wrote back to me asking why she kept getting "jumped on" for making these kinds of comment. I don't know why. I would have expected compliments. There is some kind of a fog here. Could it be the smoke from our respective burnouts?

Martin

Date: Tue Aug 25, 1992 3:30 pm PST

(ps 920825.1400)

i was quoting agre....if that wasn't clear. (i agree w/ the content, but the words as such aren't mine.)

[From Bill Powers (920825.1400)]

(á!Penni Sibun (920825) --

> why conduct elaborate deductions about yr surrounding when
>you can look and see? in particular, why maintain elaborate control
>structures when you can look and see what needs to be done? why make
>highly detailed Plans when you can improvise? why required instant
>expertise when you can improve by just keeping on doing it? why try
>figuring it out yourself when you can collaborate w/ others who have
>been there? why insist on figureing out every situation afresh when
>you can trust yor accumulated experience?

These questions, at another level, imply accepting as an explanation of behavior what I would consider a statement of the problem.

yes. but you need to ask the right questions to get the right explanations. it's *hard* to reformulate a problem in a radically diff. way, as you should know.

as representation. It sees "doing" as a process of control. Are you really suggesting that we put all that intelligence into the non-living world? If so, I'd say that you have a mammoth job of explanation ahead, if you want anyone to believe you.

the suggestion is that the intelligence is not simply and completely in the org., but in its interaction w/ the world as well.

cheers. --penni

Date: Tue Aug 25, 1992 4:07 pm PST
Subject: Re: Modeling vs. WHAT?

This discussion is gotten completely out of hand, and, like Martin, I'm rather baffled by it. I suspect that people haven't done enough homework on the other guys' stuff to justify the things they're saying about it. I rather suspect that `interactionists' have an inadequate appreciation of how feedback control systems actually work (I haven't noticed them talking about the fact the high-level systems need to respond more slowly than low-level ones, for example), but I also think that PCT is pretty shaky in the area of perception, where Chapman in particular has done useful work. But they really are very similar. Yesterday I was thinking along the lines that while Interactionism says that you can`t make breakfast following a plan, PCT goes further and says that you can't even get your hand on the spatula that way, much less your coffee-cup from the table to your lips (and that the means whereby these tasks are actually accomplished reach much further into behavioral organization than people realise).

Part of the problem, it seems to me, is that people are too eager to be believed rather than understood. It should be fine if Penni or anyone else is skeptical about the truth of PCT, as long as they actually know what they aren't believing in. But real comprehension of anything takes quite a long time to dawn - much more time than the standard dynamics of argument (and, for that manner, academic life) allow for.

I wish I had time to spend the morning trying to hose this down, but I don't, so I'll just end with a relevant quotation:

One can seldom be sure whether the silence of other scholars indicates agreement or disagreement, indifference or incomprehension.

Wilbur Maas, _Greek Metrics_.

Avery.Andrews@anu.edu.au

Date: Tue Aug 25, 1992 6:57 pm PST
Subject: Re: Modeling vs. WHAT?

I'm not convinced that PCT-ers really understand that interactionist arguments are aimed at a fairly specific group of people, the `Planning weenies', a group of people who really seem(ed?) to believe that an Agent or a critter can make its way through the world by concocting elaborate plans with its eyes closed, and then following them (still with its eyes closed). According to Chapman, for example, AI vision people actually tend to think that there needs to be an exhaustive categorization of everything in the visual field, plus precise metric information. Part of Sonja's mission is to demonstrate that this assumption is unnecessary (and, by implication wrong). The idea is not no machinery, but minimal machinery, surely scientific Right Thinking.

I think that PCT diverges from standard Interactionism by denying that much in the way of useful results will emerge from accidental outcomes of the dynamics of activity. One of Agre's examples is that if you have a stack of bowls in your cupboard, the ones you use most will tend to `drift' to the top, not because you mean for them to be there, but just as an accidental side-effect of the fact that the top of the stack is the easiest place to put a bowl you've used when you've finished with and washed it. Agre sees this as evidence of the cooperativeness of the world, while I imagine that Bill or

Rick would say that this kind of effect is of marginal importance (as would I), & most good stuff has to be controlled for. (On the other hand, setting places up so that their natural dynamics tends to support the results you want is definitely a good idea, & I suspect that Agre could make a fortune as a business consultant if he wanted to).

At any rate, we don't have to have an argument about this - analyse enough specific cases in detail, and a trend will probably emerge.

Avery.Andrews@anu.edu.au

So PCT seems cognitivist to Penni (I think) because of the emphasis on the internal representation (as reference levels) of

Date: Tue Aug 25, 1992 8:39 pm PST
Subject: The Bible is Close Too

[From Rick Marken (920825.2100)]

I feel a loose canon coming on.

Forgive me, Bill, for I care not what I do.

penni sibun (quoting Agre I think) says:

>faced w/ an empirical phenomenon to explain, our first explanatory
>recourse should be to dynamics, not to machinery.

and Powers replies:

>I don't call that an explanatory recourse; it's just describing the dynamics
>of the phenomenon (usually from the standpoint of an unspoken theory).

and Martin Taylor says

>But it is precisely the dynamics of the control hierarchy that
>describe the modelled behavior.

and I say:

Talk about generous interpretations. The Agre quote was ruling out mechanisms (which I read as "models") to explain observations. In that context, dynamics means, to me, the time course of behavioral phenomena. Bill read it exactly that way too, apparently -- dynamics refers to a DESCRIPTION, not an explanation of a phenomenon. If mechanisms are ruled out then its bye-bye control hierarchy (an invented explanatory mechanism designed to explain the dynamics of the phenomenon of control).

>it seems to me that you and she both have the same intention of simplifying
>mechanism to a point at which it can be directly tested through the dynamics.

My intention is to explain the phenomenon of control. The first order of business is to develop experimental situations where we can demonstrate the phenomenon reliably and in detail (maybe it is these details that you would call "dynamics"?). Then we see if we can build a control model to produce the results. Dynamics, per se, are not fundemental to PCT; although the world is dynamic (and, hence, so are the actions that are part of the process of controlling perceptual consequences of that world).

Martin says:

>Penni wrote back to me asking why she kept getting "jumped on" for making
>these kinds of comment. I don't know why. I would have expected
compliments.

I think Bill Powers has been MORE than diplomatic in his response to the quote by Agre. But the fact of the matter is (and since nobody likes me anyway I'll say it) that the Agre quote was a pretty concentrated heap of prattle. No one means to "jump on" anyone -- as I said, there seems to be a good deal of tolerance on the net. But if, when turning the words that are posted on the net into our own imaginings, we come up with nonsense (and that's the best I could do with the Agre quote) then we (I anyway) say it. This doesn't mean that we think Agre (or penni or anyone) is dumb. He is just WRONG (at least from my perspective). But it may be that I'm turning his words into all the wrong imaginings -- so show me what you mean with a model (mathematical or computer or mechanical); then there might be less misunderstanding.

I'm willing to believe that the Agre quote is the embodiment of the greatest wisdom of the ages. But it sounded like complete drivel to me (and Bill tried to politely explain why it did so to him as well). It would really help if penni (or anyone else who has an alternative to PCT) could describe ways we could demonstrate the meanings of the words to ourselves. If you want to know how to demonstrate the phenomenon (and model) of control to yourself, I suggest ordering my Mind Readings book for suggestions.

Martin says:

>There is some kind of a fog here. Could it be the smoke from our
>respective burnouts?

The fog seemed to be entirely in the Agre quote.

Avery Andrews (920825) says:

>This discussion is gotten completely out of hand, and, like Martin,
>I'm rather baffled by it.

I respectfully submit that there is no way that prose of the sort attributed to Agre could lead anywhere but to an "out of hand" discussion.

>I suspect that people haven't done enough homework on the other
>guys' stuff to justify the things they're saying about it.

And I suspect that this is a load of crap! I get this ALL THE TIME when I try to publish experimental tests of other people's theories. I have read more articles on other peoples theories and models than I even want to admit. I've desperately tried to find testable working versions of those models so that people could not say what you are saying. But no matter what I do, if I create a situation that exposes the failures of the model, I am quickly told that I don't understand the model or that the model doesn't really work that way. Either I'm really stupid (a serious possibility -- just ask my kids) or there is a very cute shell game going on out there in ScienceLand that I hadn't intended to play.

>I rather suspect that `interactionists' have an inadequate
>appreciation of how feedback control systems actually work.

Actually, what they have is no idea what feedback control system DO. They know zip about control so they obviously have no interest in figuring out how control systems work. Many of these people (ai types mainly) have no apparent interest in science, for that matter. They just like to play with models (or philosophical abstractions) and have no idea that there is a REAL phenomenon out there (called control) BEGGING for an explanation.

>but I also think that PCT is pretty shaky in the area of perception,
>where Chapman in particular has done useful work.

What in the world is "shaky" about the PCT view of perception???? What perceptual processing has Chapman put in his models that is so "useful". The Agre and Chapman article that I read reflected one of the most naive and unsophisticated views of perception I've ever seen coming out of grown-ups. Have you read Behavior: The control of perception?? Have you ever seen a behavioral model with a more sophisticated, deeper, functional understanding of perception (maybe you mean the brilliant perceptual model of stimulus sampling theory? dynamical systems theory? what???).

>But they really are very similar.

And the biblical description of genesis is "very similar" to the cosmological/evolutionary description of genesis.

Doesn't anyone read my rants? Remember. All these people with theories that are "very similar" to PCT forget that little part about how organisms CONTROL and that what they control is PERCEPTION. Did anyone read my "Blind men" paper. Did anyone read Bill's Psych Review article. Should I write it in RED. How about repeating it:

Behavior is controlled perception;
Behavior is controlled perception;
Behavior is controlled perception;
Behavior is controlled perception.

That doesn't make it true, but I hope it does make it NOTICEABLE. Now ask all these people with similar theories if the above collection of phrases strikes them as being a good description of their model of behavior.

>Part of the problem, it seems to me, is that people are too eager to be
>believed rather than understood.

I don't want to be believed. I want to be understood. That's why we develop all these demos and experiments and models. No need for belief; just LOOK and TEST.

>It should be fine if Penni or anyone else is skeptical about the truth of PCT.

Absolutely. And she is free to maintain her skepticism. I'm not objecting to her skepticism. I'm only saying that the stuff quoted in Agre is drivel. But I am willing to be convinced that I am wrong. Just show me how Agre's theory explains the one phenomenon that I know exists (because I can demonstrate it) and that seems to characterize everything that we call human behavior -- CONTROL.

>I wish I had time to spend the morning trying to hose this down,
>but I don't, so I'll just end with a relevant quotation:

>One can seldom be sure whether the silence of other scholars
>indicates agreement or disagreement, indifference or incomprehension.

> Wilbur Maas, _Greek Metrics_.

I'll give you odds that its the last three.

There. Now Bill can play GOOD COP.

Regards

The loose canon

(Yes, I know how to spell cannon)

Richard S. Marken

Date: Tue Aug 25, 1992 10:27 pm PST
Subject: Re: The Bible is Close Too

[from Avery Andrews 920826.1615]

Rick Marken (920825.2100)

> The Agre >quote was ruling out mechanisms (which I read as
>"models") to explain observations.

No, it wasn't. But it probably has to be read in context to be understood properly. It was advocating that mechanism by reduced to a minimum (but a *sufficient* minimum, which puts C&A in a different basket than, say, Bickhard, who seems to have no sense at all of what `sufficient mechanism' might look like). He wants people to get prizes for eliminating unnecessary mechanism.

>My intention is to explain the phenomenon of control.

Fine, but it is also necessary to get people to suspect that control is central to ordinary human activity, and that they don't already understand it properly. This, I think, requires having a large library of examples that will `grab' people, & picking the ones that will appeal to people from particular backgrounds (given the fragmentation of the human `sciences' there won't be any `one size fits all' solution. Somehow, for some reason, that fact that you need feedback to point at things or stand upright doesn't seem to get to AI-ers, I think because they don't seem `cognitive' enough. The right examples, I think, have a capacity to burn holes in people's minds that mere arguments don't (academics are trained to win arguments, and people from Harvard and MIT base their egos on the the idea that they're (almost, if they're humble) the smartest person they know). I'm hoping that the steering example will prove effective, but there must be lots more.

>Actually, what they have is no idea what feedback control system DO.
>They know zip about control so they obviously have no interest in
>figuring out how control systems work.

I *think* that this is true. So what's needed is to get them interested.

>any of these people (ai types mainly) have no apparent interest in science,
>for that matter.

Definitely false of the Interactive AI crowd, from what I've seen.

>What in the world is "shaky" about the PCT view of perception???? What
>perceptual processing has Chapman put in his models that is so "useful".

Well, could you build a system that could look at a video screen & process what it sees there well enough to play the game? Chapman can't really do that either, but Sonja embodies a lot of work on the problem. Interestingly, in his thesis proposal, he said that Sonja' perceptual system would be like Pengi's, but her central processor would be much more complicated. Interestingly, the reverse proved to be true: `deciding what to do is easy when you know what's in front of you'. The main reason I didn't become a PCT-er at age 20 or so was an inabilty to get anywhere on understanding how high-level perceptual functions work - e.g. not a clue about how language is actually processed, for example.

Now, I think it's time to look again, partly because there has been real progress on the perceptual front, also because you & Bill have gotten me to begin to suspect that maybe the problem of unknown mechanisms for perception is not as serious as I thought it was (if you can demonstrate that the critter is controlling for X, then you know it contains an X-detector, even if you don't know how to build one. And you know that you ought to look for one).

Avery.Andrews@anu.edu.au

Date: Tue Aug 25, 1992 10:54 pm PST

[From Hank Folson (920826)]

(Chris Malcolm 920825)

>Clearly these apparent paradoxes arise because the "folk" theories of
>how control is learnt and executed are wrong.

Is there really a paradox here? As you described, the only way to get a motorcycle/bicycle to turn right is to get it leaning to the right before initiating the final steering actions. The only reason we make a big deal about counter-steering is because i t is not obvious that we must do it and do do it. IsnUt the use of the term "counter-steering" a stimulus-response sort of description?

>What is required is a theory of control which can be applied to the
>phenomena of counter-steering in which the paradoxes become predicted.

How about this, Chris: One thing that I am coming to understand about control theory is that purpose is all that counts. Method is secondary, just a means to the intended end. The motorcycle/bicycle steering scenario is a good example of this, I think.

The first time someone without a physics degree or an experienced tutor tries to ride a bike/motorcycle, they will steer right when they want to turn right. The desired perception is turning right, and turning the handlebars to the right seems like a good idea at the time. The result is a fall or a moment of panic and wobbling with no turn. The purpose/goal of turning right remains, however. Steering right didn't work, so I suspect that our control systems simply try another controlling action AT RANDOM.

Handlebars only go right or left, so the next controlling action is to turn the bars to the left. NOTE THAT THE SYSTEM HAS NO INKLING OF WHY THIS SHOULD WORK ANY BETTER THAN TURNING THE BARS TO THE RIGHT JUST DID. The bike now falls to the right side in r esponse to God and Newton. Whatever level of our hierarchy controls for not falling quickly takes over. The most recent output was to turn the handlebars to the left. So now the system corrects, and turns the bars to the right. This is the only controllin g action option available, fortunately. (In other words there are not so many degrees of freedom that all options could not be tried before you fall to the ground.)

Now the bike has velocity, and is leaning to the right, and the bars are turned to the right. With the centrifugal (or is it centripedal?) force, and the human control system hopefully in a stable range, there is no fall, just a beautifully controlled rig ht turn EVEN THOUGH THE CONTROL SYSTEM DOES NOT UNDERSTAND WHY.

There is not much in the way of planned output in my perception of how a control system functions. The initial controlling action chosen may be selected for some reason (e.g. If I want to turn right, it makes sense to turn the bars to the right, based on what is stored in the mind from past experiences, such as steering a tricycle.), but if that first choice does not produce the desired input as the process goes around the control loop, or if there is no reason to choose a particular action, the control system will try one at random. And if that doesn't work, it will continue to control by trying another action. One thing I like about this suggested interpretation is that a living control system can survive a great many complex and/or new situations, with out requiring a lot of computing power, and can react very fast. Perhaps those lucky enough to randomly pick the best controlling actions first survive better than others!

The bicycle steering example is, to me, a good example how a control system can, by random actions, handle something that its sensors and internal logic would never figure out. The way E. coli tumbles and takes off in a random direction when controlling for a new food source is a lower level example.

In the going right by steering left scenario, I did not mention another controlling action the system has available: putting a foot down. This is perfectly valid, and many beginners probably do this. This takes care of the error signal of falling down, bu t doesn't do anything about turning right. Those who put their foot down will probably take longer to learn how to ride, because their control systems will not get the opportunity to complete the steering left to go right strategy.

Another thought to ponder: The system has just learned by trial and error how to turn right. How will the system respond to the next new goal which would be to turn left? Will the rider try steering left to go left, or will he immediately steer right to go left? Is this a prediction of your paradox, Chris?

Hank Folson, Henry James Bicycles, Inc.
704 Elvira Avenue, Redondo Beach, CA 90277
310-540-1552 (Day & Evening) MCI MAIL: 509-6370
Internet: 5096370@MCIMAIL.COM

Date: Wed Aug 26, 1992 4:58 am PST

From Greg Williams (920826)

Rick Marken Says:

>Does anyone know of any written critique of Behaviorism which explicitly
>cites the following problem (that Gary mentioned some time ago) :if
>human behavior is controlled by the environment (as claimed by Behaviorists)
>then Behaviorists themselves should not be able to exert control over people
>(as they say that they can -- and should) because they are under control?
>What I want is a critique that points out that control cannot be exerted
>by agents that are controlled. I have never seen any detailed critique
>of Behaviorism from this perspective.

I've not seen any either -- and I've spent a lot of time investigating behaviorist philosophy -- probably because the critics are wise enough to realize that the behaviorists can wheedle out of the charge. When backed against the wall, they are perfectly willing to admit that their own environmental and genetic histories have determined their own present-time "control" of other organisms, where "control" simply means making certain types of rearrangements of the environment of the other "controlled" organisms. And they would even admit that (sometimes) the organisms being "controlled" rearrange aspects of their (the behaviorists') environments, resulting in a modicum of mutual simultaneous control. They realize that an organism's environment can be altered by other organisms -- but they claim that ultimate "control" of any organism is exercised by its entire environmental and genetic history, so that, for example, an experimenter's life history "made" him/her starve rats and then "control" what their actions to get food. Any ability to alter another organism's environment can result, potentially, in "control" of that organism. So the behaviorist's college professors helped to "make" him/her (later) starve rats, etc. And the professors' families helped to "make" them (later) "make" behaviorists (later) starve rats, etc., etc., etc.

The real problems in this self-consistent scheme are (1) not having sufficiently detailed (generative) models which can predict the LIMITS of "control" (i.e., starving THOSE rats results in WHAT DEGREE OF "control"?) and (2) invoking "history" as a catch-all explanatory notion (and then concocting "just-so" Whiggish stories, as do some evolutionary theorists). But the behaviorists say that (1) their models are good enough for their purposes and (2) that invoking history is as good as we can do right now, given the state-of-the-art in neurophysiology. (Sometimes when I read speculative musings on the net, I'm rather sympathetic to the latter claim!)

Greg

Date: Wed Aug 26, 1992 7:24 am PST
Subject: Re: The Bible is Close Too

[Martin Taylor 920826 11:00]
(Rick Marken 920825.2100 canon fire)

> Dynamics, per se, are not fundemental to PCT; although the world is dynamic
>(and, hence, so are the actions that are part of the process of controlling
>perceptual consequences of that world).

In this quote is the heart of the resistance to using information theory and the results of psychophysical studies. Rick is controlling for the perception that dynamics is irrelevant to the operation of a control system, whereas I believe that the control system is defined by its dynamic behaviour. Not everything can be relegated to the result at infinite time, as the stability equations often trotted out assume. It matters that the control system respond fast enough to changing events in the world (disturbances). It matters (sometimes) whether there are overshoots in the control function.

Why is there a conceptual problem with understanding the importance of the differing numbers of degrees of freedom in different parts of the hierarchy? Because they depend on the dynamics, and dynamics is predetermined to be necessarily unimportant. A static analysis is simply not good enough.

I don't know why Rick, in particular, is controlling for this reference, because it contradicts his own presentation (I've forgotten the title of the paper for the moment) pointing out the congruity between the rates at which certain actions can be performed and the rates at which the relevant perceptions can be performed. If that isn't a consequence of dynamics, what is? Why are the so-called "slowing factors" included as a naive method of ensuring stability in the hierarchy simulations if not for reasons of dynamics?

Dynamics is central to the importance of PCT as a psychological theory. As a static theory it only presents a trivial truth that has some amusing consequences. You'll never persuade the world of its beauty by repetitive ranting based on that trivial truth.

>Forgive me, Bill, for I care not what I do.

Your resistance to disturbances leads the outside observer to doubt the truth of this claim.

Martin

Date: Wed Aug 26, 1992 9:30 am PST
Subject: What's important

[From Rick Marken (920826.1000)]

>From Greg Williams (920826)

Thanks!!! Very helpful.

Avery Andrews (920826.1615)

>Somehow, for some reason, that fact that you need feedback to point
>at things or stand upright doesn't seem to get to AI-ers

The "Blind men" paper shows that feedback is THERE whether you need it or not. Organisms are ipso facto locked in a loop and behaviorally it is clear that the sense of the feedback in this loop is negative. Who cares what ai-ers think? I'm not much interested in the Pope's ideas about population growth either.

> The right examples, I think, have
>a capacity to burn holes in people's minds that mere arguments don't

We've been demoing this stuff for YEARS. Obviously using the WRONG examples. I'm pretty convinced that there are no right examples (any more than there are the right stimuli to make people do stuff). The only thing that can be right is the willingness to learn -- and I haven't seen much of that amongst ai types. In fact, the usual approach is to explain to us why our demos and models are not (or don't need to be) doing what they are doing -- controlling perception.

>Well, could you build a system that could look at a video screen &
>process what it sees there well enough to play the game?

Sure. Believe me, we can design perceptual systems with the best of them; and if we can't, we are happy to learn how they do it and incorporate those results into our models, when necessary. I have always said that the main substantive problem for PCT behavior modelling will be learning how to build functions that perceive some of the more complex variables that we obviously control. The whole idea of PCT is that behavior IS perception in action -- controlled perception. It's all about perception. So if anybody is building great perceptual models, that is not an argument agaist PCT and for the approach to behavior of the perceptual modeller -- it just shows that there is information about perception that we can use.

Suppose, for example, that we discover that "honesty" is something people control. We certainly can't build an "honesty detector" now (or in the fore-seeable future). Would such a situation make PCT less plausible? NOT. Suppose the advocates of some other theory could build an honesty detector. Would that make the other theory more plausible than PCT. NOT. But it would certainly suggest that we should look at their honesty detector algorithm for use in our PCT model.

PCT modellers are well aware of the problems of modelling perceptual functions and we TRY to keep up with the state of the art. So I think we are at least at the level of a Chapman and Agre, unquestionably at the level of a Beer or Brooks.

Martin Taylor (920826 11:00)

> Rick is controlling for the
>perception that dynamics is irrelevant to the operation of a control system,
>whereas I believe that the control system is defined by its dynamic
behaviour.

Dynamics are NOT irrelevant to the operation of control systems; if the dynamics are wrong, there is NO control. But I'm not particularly interested in modelling the dynamics, per se. Once the system works (so that it behaves (dynamically) just like the real system) then I'm happy. I'm MUCH more interested in WHAT is controlled (what perceptions), WHY its controlled (higher level goals?) and HOW its controlled (use of lower level perecptions?). Dynamics "fall out" of the operation of the control model. They certainly have to be there (and appropriately -- witness the discovery of the need for a lag to account for the marken effect) but it sounds like you want to model the dynamics itself. That's fine with me (they have been doing that in human factors for years -- and missing completely the MAIN point, that these systems are controlling perceptual variables relative to internally specified references for those variables).

>Not everything can be relegated to the result at infinite time, as the
>stability equations often trotted out assume.

I agree. But once you know that there is a dynamically stable solution to the control equations, then you can present the algebraic results to clarify the main points of control -- control of perception, resistance to disturbance. You can summarize the main properties of a lever without giving the differential equations for it, can't you??

>Why is there a conceptual problem with understanding the importance of the
>differing numbers of degrees of freedom in different parts of the hierarchy?

Because I don't know what phenomenon this relates to. I am interested in modelling human behavior; that's what I do well. I can't do all these fancy mathematical analyses of properties of the model itself. If you want to do that then go ahead. You may find out something important about the model and that would be great. But I'm more interested in the phenomenon of control. So if your degrees of freedom analyses of the model predict something about the phenomenon of control, then let me know and I'll test it.

>Dynamics is central to the importance of PCT as a psychological theory. As
>a static theory it only presents a trivial truth that has some amusing
>consequences. You'll never persuade the world of its beauty by repetitive
>ranting based on that trivial truth.

I don't think I could possibly disagree with you more. I presume the trivial truths revealed by the static equations are

p = r and o = -kd

Perception (p) is made to match an internally specified reference (r) and this is accomplished by producing outputs (o) that cancel the effects of disturbances (d) on p.

OK. There's the trivial truth.[If it's so trivial, why don't psychologists TAKE IT FOR GRANTED AND ACT AS THOUGH THEY UNDERSTOOD IT (like by not doing research in the stupid way it is currently done)]. Now, please explain the non-trivial truth that is revealed by dynamic analysis.

>>Forgive me, Bill, for I care not what I do.

>Your resistance to disturbances leads the outside observer to doubt the
>truth of this claim.

That's the fact, Jack.

Regards   Rick

Date: Wed Aug 26, 1992 11:23 am PST
Subject: Re: What's important

[Martin Taylor 920826 14:45]
(Rick Marken 920826.1000)

>>Why is there a conceptual problem with understanding the importance of the
>>differing numbers of degrees of freedom in different parts of the hierarchy?
>
>Because I don't know what phenomenon this relates to. I am interested in
>modelling human behavior; that's what I do well. I can't do all these fancy
>mathematical analyses of properties of the mdoel itself. If you want to
>do that then go ahead. You may find out something important about the
>model and that would be great. But I'm more interested in the phenomenon
>of control. So if your degrees of freedom analyses of the model predict
>something about the phenomenon of control, then let me know and I'll test it.
>

Well, I thought I had done exactly that, which is why the plaint above. I demonstrated the absolute need for some means to pass control from one ECS to another (or, equivalently, a means to change the perceptual input function of controlling ECSs). If that isn't an important point about the model, I don't know what is. It deals all at the same time with attentional focus and with wide-open alerting systems. And if you don't understand it simply because the relation of degrees of freedom to system dynamics is unimportant, then you miss something critical in PCT (whether H or not).

>Perception (p) is made to match an internally specified reference (r) and
>this is accomplished by producing outputs (o) that cancel the effects of
>disturbances (d) on p.
>

>OK. There's the trivail truth.[If it's so trivial, why don't psychologists
>TAKE IT FOR GRANTED AND ACT AS THOUGH THEY UNDERSTOOD IT (like by not doing
>research in the stupid way it is currently done)].
>

I can't answer the first part of the question, because to my mind it is a truth that is gloriously self-evident once someone points it out. It is the initial discovery of that truth that is the sign of genius. As for the second part, I can answer from personal experience. Even when you acknowledge the fundamental truth, its consequences and power are far from self-evident. Even Bill says he's still learning about that aspect, and for myself I have seen only that there is glory there, not its shape.

Martin

Date: Wed Aug 26, 1992 11:58 am PST
Subject: wet mind

I am enjoying reading a book now and thought that other list members might like to know about it.

Title: Wet Mind
Author: Stephen M. Kosslyn & Oliver Koenig
Publisher: The Free Press, 1992

If anyone has read this book and would like to talk about it in relationship to HPCT, I am game.

David Goldstein

Date: Wed Aug 26, 1992 1:45 pm PST
Subject: This and That

from Ed Ford (920826:1400)

Last week I gave my first workshop on the Outcome Based Parenting program I created which includes a heavy amount of PCT. The PCT went extremely well. The participants saw and accepted the many connections between PCT and human thinking and actions. More on that later. When I returned, I found my computer had been permanently zapped by lightening. One son living at home lost an amplifier and another son several miles from my home lost a new TV set. Some blame it on that crazy weather (notice I didn't say that crazy man) from California. At any rate, after making the purchase of a new computer and having (aren't I lucky) my son-in-law get me back in operation yesterday (my first back up disc didn't work, but the second did), I found a lot of mail.

Mark, give me some time and I'll reply to your thoughts on dealing with the past.

Rick, I'm a Republican....more on your convention comments later.

Bill, my perception of the word belief (noun) is tied directly to what I have found to be true and/or working through dealing with others, thinking through a problem, and/or evidence. That's what I mean by belief. Faith (noun) is the belief in something about which you have no proof or at least insufficient evidence. When it comes to a person's religion, it can be a faith or a belief. Mine happens to be a belief.

Ed Ford

Date: Wed Aug 26, 1992 3:42 pm PST
Subject: Re: Behaviorist critique

[Martin Taylor 920826 1900]
Rick Marken (Quoted by Greg Williams 920826, initial date not immediately accessible)

>Does anyone know of any written critique of Behaviorism which explicitly
>cites the following problem (that Gary mentioned some time ago) :if
>human behavior is controlled by the environment (as claimed by Behaviorists)
>then Behaviorists themselves should not be able to exert control over people
>(as they say that they can -- and should) because they are under control?
>What I want is a critique that points out that control cannot be exerted
>by agents that are controlled. I have never seen any detailed critique
>of Behaviorism from this perspective.

I'm afraid I don't see the paradox. What is the problem with controlled agents themselves controlling others? Let us take the S-R position that some pattern in the environment determines the behaviour of an agent (I assume this is what "control" means in that context). Then why should the actions of that agent not create a pattern in the environment that determines the behaviour of another (sub)agent?

It seems to me analogous to the situation of a high-level ECS controlling its percept by way of (unknown to it, of course) setting references that lower ECSs use in controlling their percepts. Setting a reference level is the closest behavioural thing in PCT to the S-R idea of control, I think. Anyway, whether you buy this analogy or not, I can't see the paradox when the controlled controller is seen strictly from an S-R viewpoint.

(Incidentally, I'm still very uncomfortable with the conflation of "behaviorist" and "S-R viewpoint." As I like to use the word, any psychologist who is not a behaviourist should be considered a mystic. PCTers are not mystics. They observe what happens in the world and model it using plausible, possible mechanisms that are as similar as possible to what is known from real observations of real people and physically described machines. PCT is a thoroughgoing behavioural theory, in my view. And I think you know by now that I do have a slight inkling of what PCT is all about.)

Martin

Date: Wed Aug 26, 1992 4:01 pm PST
Subject: Re: Modeling vs. WHAT?

(ps 920826)

From: Avery Andrews <andaling@durras.anu.edu.au>

I'm not convinced that PCT-ers really understand that interactionist arguments are aimed at a fairly specific group of people, the `Planning weenies', a group of people who really seem(ed?) to believe that an Agent or a critter can make its way through the world by concocting elaborate plans with its eyes closed, and then following them (still with its eyes closed).

i want to emend this just a bit. most of the authors we've been discussing are coming mainly out of (and addressing) ai. however, there's definitely collaboration w/ other fields, particularly philosophy (cf. preston), anthropology (ethnomethodology/conversation analysis), and (some) psychologists. the interactionist argument is in general deemphasizes individuals and emphasises their interactions.

it's an intellectual movement that is taking hold and still growing and gradually being assimilated, and it has its merits and demerits like any other. i think it's not going to go away, at least insofar as it gets assimilated. but it's not the ultimate truth, and more than anything else (incl pct) is.

cheers. --penni

people actually tend to think that there needs to be an exhaustive categorization of everything in the visual field, plus precise metric information. Part of Sonja's mission is to demonstrate that this assumption is unnecessary (and, by implication wrong). The idea is not no machinery, but minimal machinery, surely scientific Right Thinking.

I think that PCT diverges from standard Interactionism by denying that much in the way of useful results will emerge from accidental outcomes of the dynamics of activity. One of Agre's examples is that if you have a stack of bowls in your cupboard, the ones you use most will tend to `drift' to the top, not because you mean for them to be there, but just as an accidental side-effect of the fact that the top of the stack is the easiest place to put a bowl you've used when you've finished with and washed it. Agre sees this as evidence of the cooperativeness of the world, while I imagine that Bill or Rick would say that this kind of effect is of marginal importance (as would I), & most good stuff has to be controlled for. (On the other hand, setting places up so that their natural dynamics tends to support the results you want is definitely a good idea, & I suspect that Agre could make a fortune as a business consultant if he wanted to).

At any rate, we don't have to have an argument about this - analyse enough specific cases in detail, and a trend will probably emerge.

Avery.Andrews@anu.edu.au

So PCT seems cognitivist to Penni
(I think) because of the emphasis on the internal representation (as reference levels) of

Date: Wed Aug 26, 1992 4:02 pm PST
Subject: Outcome Based Parenting program

Ed,

Can you tell us about your Outcome Based Parenting program?

Thanks, David Goldstein

Date: Wed Aug 26, 1992 5:04 pm PST
Subject: Modelling vs. What

Re Rick's complaints on the Chapman & Agre papers - I found them rather obscure too, so I borrowed the books from Penni & then got my own copies. I think it's pretty normal that papers are enigmatic unless you already understand the background pretty well. I find that AI papers tend to leave out too much of the substance of what's being done for me to follow them, but I imagine that people with the right training would not have this problem. Books present more of the background, & are more detailed about the substance.

Moving on to a more specific point, Bill's arm demo (the version I have at least) has to assume that the visual system can track the fingertip & the target, whereas a major aspect of Sonja is an visual system that can to some extent actually do this (though he ducks problems of early vision).

I actually think that together, Bill and Chapman could build a robot that could play real video games, at least if you wheeled it into the arcade and plunked it down in the one it was supposed to play. But I don't think either could do it alone.

And, on PCT presentation, it would be nice to have more in the way of concrete examples of 2nd order control systems doing useful things. I think I understand how the vertical angle control inthe arm demo works, but that's actually the only example I know, & it would be nice to have more. Maybe they're in the stuff I've ordered from Pat&Greg, but still it might be a good idea to have it available online.

Avery.Andrews@anu.edu.au

Date: Wed Aug 26, 1992 8:23 pm PST
Subject: Misunderstandings & miscellany

[From Bill Powers (920826.2000)]

Spent a fine morning at Hovenweep ("Deserted Place" in Ute) National Monument just across the border into Utah. Variation on Chaco Canyon construction, but still the same idea: pile up big and little rocks with adobe mud between them. It struck me that having a Plan for stacking up rocks would have been futile when the Anasazi builders put up one of those round, square, or D-shaped towers that abound at Hovenweep. How can you make a Plan of Action when you don't know exactly how one stone is going to fit on another, or how many thin stones you'll need to make the next course level? The only Plan you can have is for how the thing is supposed to look when it's done, and they clearly had that. When I look at those T-shaped doors, the little windows between rooms, the slanted peepholes, the straight corners on some buildings and nicely rounded walls on others, I get the sense of seeing into minds nine centuries dead. They have left their reference images behind for us to see, in stone.

Avery Andrews (920825) --

Avery, you are becoming the voice of reason on this net. Keep it up. You will get us back on track.

I have received your C program. It is easy to read and understand, which to me is good programming. I'll try it tomorrow -- no reason why it shouldn't work. Only a few very small changes ought to make it work even better. Back to you later with some ideas about how to tie Astro to language.

Penni Sibun (previous posts) --

From all the comments by others, I suspect that I misunderstood your last post in most of the possible ways. My apologies. As Avery said, we are ALL arguing against somebody else's point of view without really understanding it. You keep your temper remarkably well, which should be a lesson to all of us.

Just to clear the air, I interpreted your rhetorical questions as if they were addressed to PCTers, with the "why" part a (misplaced) criticism of PCT and the "when" part your suggested remedy. Now I realize that this was a quote and not your words at all, and that the target was the old-fashioned AI weenies. So we really have a three-way collision going on here (4 counting interactionists, if I understand correctly). It will help if we can keep straight what aspect of whose ideas is being discussed, and why. I'm an old, old, man, Penni, so you have to forgive me (heh, heh).

With Avery's explanations, I can certainly go along with the idea of looking for the minimum necessary machinery. My philosophy of modeling exactly.

Hank Folson (920826) --

Very nice elaboration on the riding-the-bicycle thread.

Greg Williams (920826) --

There's one problem with the "self-consistent" behaviorist scheme that you didn't mention: the reference signal. When a behaviorist trains a rat to do something, what the rat is to do is already known to the experimenter. As I've said before, a behaviorist who wants to demonstrate the reality of stimulus control of behavior must state BEFORE the demonstration what behavior is going to be brought about. In other words, the experimenter must describe what will be perceived when the process has finished.

Skinner himself discussed shaping in those terms. He said that the experimenter should simply keep in mind the final pattern that is to be performed by the animal, and reinforce any little move that looks like a step in the right direction. He emphasized that there's no standard procedure here; he said in so many words that the experimenter must be prepared to improvise. So he was describing the behavior of a control system, but trying to explain it without using any CT concepts -- avoiding particularly any mention of intention or purpose. Without using those concepts, of course, he would have been unable to explain the actions of the experimenter.

When Skinner spoke about an experimenter controlling an animal's behavior, his language always contained some way of saying that the experimenter desired, specified, or intended that a certain behavior come into being. In talking about the environment controlling behavior, of course, he did not use such terms. He didn't say that the environment had in mind a particular behavior of the organism, and improvised rules and happenings that reinforced any little move toward that behavior. In speaking of "organisms" (meaning non-experimenters), he also had to avoid anything sounding like a purpose, such as knowledge in advance of the outcome or a desire for a particular outcome. Only the experimenter was allowed to behave in a purposive way, knowing what was to happen and doing whatever was required to make it happen. Therefore he attributed true control only to behavioristic experimenters, and not to subjects or to environments.

When he use the term "control" with respect to non-scientific organisms or environments, he did not mean goal-directed error- correcting actions like those the experimenter used. He meant "influence" or "determine" or "cause" -- purposeless cause and purposeless effect. In that context, "control" was a code-word, a euphemism. He was doing as he often did: he used a common word that to most people implies intention or purpose, but with the understanding that only his meaning was allowed. Something like the way certain politicians are now talking about "family values."

So I dispute use of the term "self-consistent" in reference to behavioristic notions of control.

Martin Taylor (920826.1100) --

RE: dynamics. I was interpreting Penni's reference to putting a study of dynamics first as meaning that we should observe only the dynamics of interactions as an external observer would see them. I think Rick may have put the same interpretion on this passage. In fact, Penni is right. Before we propose any model, we should first observe the behaving system in its environment, and note the dynamics of the interaction (by which I don't mean "psychodynamics," which involves heavy layers of theory, but simply the temporal aspects of the physical organism-environment relationship). This, I still maintain, amounts to laying out the phenomenon that needs an explanation.

Then, of course, a model is required to explain the observed dynamics. This model will have to have certain internal dynamical properties of its own in order to show the same external dynamics as the real system interacting with a real environment. So if I have misunderstood what Penni was saying, at least my misunderstanding should be clear now.

RE: degrees of freedom.

>I demonstrated the absolute need for some means to pass control from
>one ECS to another (or, equivalently, a means to change theperceptual
>input function of controlling ECSs).

I think we will all continue to be frustrated on this subject until we come up with experiments that demonstrate this "absolute need." You are talking about a HYPOTHETICALLY absolute need, conditional on certain premises being true. What we need is a demonstration with a live human subject in which a phenomenon occurs that can only be explained by this organization in a model. It should be something simple enough to do on a PC, so that the modelers on the net can do the experiments, program the model, and test the match. We can stipulate (as Avery pointed out) the existence of perceptual devices as required to provide perceptual signals representing whatever aspects of the environment are needed. The actions required should be simple motor behaviors, to accomodate mouses and joysticks, but these behaviors can be translated into any special effects on the screen desired, simulating other kinds of behaviors (such as cranking, pressing buttons, uttering words, and so on).

My proposal that a given ECS can be applied to different lower-level perceptions at different times is somewhat similar to yours. Mine is no better founded in experiment than yours. I believe we have approached the limits of what can be accomplished by passing words back and forth. We need to devise critical experiments and carry them out. Otherwise we might as well be arguing about angels and pinheads.

Ed Ford (920826) --

Congrats on your new computer. And most particularly, congrats doubled for having not only one backup disc, but two.

I have now read "Outcomes Drive the System," and I can see why you think that ODDM (Outcomes Driven Development Model) is ripe for a merger with control theory.

To all and sundry. ODDM is a method for organizing educational systems according to some basic assumptions. The author of this booklet (unnamed but possibly John Champlin) lists them this way:

1. What I want determines what I do.

2. What I do must reconcile and align with the most current, valid knowledge.

3. All actions and all decisions are aligned and intentional.

.... [skipping a few rather vague generalities]

6. There is no condition [or] feature of behavior which cannot be intenionally altered in such a way as to insure success.

7. Change is a holistic activity. Single parts cannot be altered without significant impact in all of the organization.

8. Any change must be in response to an established and accepted need.

...[and another generality].

Later:

A need is created by

(a) a discrepancy between what we want and the condition presently existing.

....

(c) a difference between what we claim to be doing versus what one actually observes.

>Rick, I'm a Republican....more on your convention comments later.

No, Ed, you are not a Republican. You are a control theorist. Being a control theorist, you can find aspects of the Republican point of view that are consistent with human nature, and others that are not. You can do the same with the Democrat point of view, or Sun Yung Moon's. A theory of human nature is contained in every political point of view -- but is in important ways unfounded, because political points of view are not based on any underlying principles that can be tested. I put it to you that PCT and HPCT transcend political points of view, as well as many others. Except for Rick. Yukayukayuk.

Belief: Each of us has criteria for what constitutes enough evidence to justify belief (i.e., acting and thinking as if something is true). If the true foundation of belief is evidence, then new evidence or a reconsideration of old evidence can easily change a belief, because in that case a belief is the result of evidence, not the reason for accepting evidence. I define faith the other way around: faith exists when the belief comes first, and the evidence is selected in any way required to prove the belief true. So faith is preserved in spite of the facts, while belief is generated because of the facts.

Of course there are several layers to this concept: what do you believe to be the facts? And why?

Best to all, Bill P.

Date: Wed Aug 26, 1992 8:51 pm PST
Subject: Gathering Interactionism

[Avery Andrews 920827.1428]

On the general theme of Interactionism, the GURU parameters of CROWD (= Gatherings) produce a very nice Interactionist result: people to tend to gather in arcs around gurus at gatherings, and a cognitivist might suppose that this happens because the individuals are trying to make their collective configuration match an `arc around the guru' schema. But Bill's model shows that the phenomenon could be a consequence of the dynamics you get from certain reference levels for proximity and distance.

And in the absence of real evidence that the arc configurations are being controlled for, the dynamics explanation is clearly superior, on grounds of parsimony.

Avery.Andrews@anu.edu.au

Date: Wed Aug 26, 1992 11:59 pm PST

Subject: Re: Misunderstandings & miscellany

[Martin Taylor 920827 0020] (Bill Powers 920826.2000)

Talking about my claim to have demonstrated a novel and important fact about PCT through consideration of dynamics and the associated degrees of freedom...

> What we need is a demonstration with a
>live human subject in which a phenomenon occurs that can only be
>explained by this organization in a model. It should be something
>simple enough to do on a PC, so that the modelers on the net can do
>the experiments, program the model, and test the match. We can
>stipulate (as Avery pointed out) the existence of perceptual devices
>as required to provide perceptual signals representing whatever
>aspects of the environment are needed.

I doubt you could ever find a phenomenon that can be described by only one organization. If you want an anecdotal set of facts that fit the proposed organization, then you need only look at someone's head turn to face the source of an unexpected noise, or the flick of the eyes toward a distant light flash.

The argument goes the other way: if there are more sensory df than motor df (either instantaneously or over time) and if a system is a living one (i.e. has been built by Darwinian evolution), and if PCT is valid for living systems, THEN there must be passive alerting systems that can direct control among the possible percepts permitted by the sensory systems.

I make no claim for any specific organization of how these alerting systems might work. I only claim that different percepts must come under control at different times, and that there must be some mechanism for shifting control to those percepts it is most important to control at any particular moment. The instigator of that change is what I call the alerting system, which might well not be a separate system from the control hierarchy, but which must work passively and effectively in parallel across all the percepts that might be controlled (or at least across those not currently being controlled). Your proposed organization might fulfil these criteria. I don't know.

There really isn't any experimental demonstration that I can think of, because you can't unbuild an evolved living system, and if you deliberately build a simulation to have alerting systems that can affect what is controlled, then you are testing a tautology. I guess you might set up (or attempt to set up) a negative demonstration--a system that had far more sensory inputs than motor outputs, that could control appropriately many percepts at several levels, and that had no passive parallel alerting system. Success at such a demonstration would show that such an alerting system is not necessary, but I think it would be necessary to provide an equivalent (e.g. an autonomous scanning system checking on the errors of currently non-controlling ECSs).

This basic condition for any such simulation is that the sensory system of the simulated controller be much broader-band than the motor system, and that the environment be rich enough to justify the sensory bandwidth. What this means is that the experimenter-simulator must set up lots of CEVs that the simulated organism could learn or be constructed to control (depending on whether we wanted to bring an experiment on reorganization into this). It may be that it would be easiest to do this in the natural world. But perhaps one could work with a very restricted world, using a 2 df motor system and, say a 6df sensory system of bandwidth higher than the motor system by an order of magnitude. That would give 30 times as many sensory df per second as motor df per second. Even this, I think, would be hard to mate with a simulated world, especially when the objective is to show that the organism can control a considerable variety of percepts, for each of which the environment must provide the data, at several justifiable levels.

Maybe more experienced modellers than me could see a way to try such a negative demonstration. Maybe someone else can think of a way to do a positive demo or experiment. I think of it as a fact on the same level as that you can't keep a cursor on a moving target unless you can see at least the deviation of the cursor from the target.

Martin

Date: Thu Aug 27, 1992 5:45 am PST

From Greg Williams (920927)

Before it gets damp here from Andrew's remains:

>Martin Taylor 920826 1900

>I'm afraid I don't see the paradox. What is the problem with controlled
>agents themselves controlling others? Let us take the S-R position
>that some pattern in the environment determines the behaviour of an agent
>(I assume this is what "control" means in that context). Then why should
>the actions of that agent not create a pattern in the environment that
>determines the behaviour of another (sub)agent?

The "paradox" arises for PCTers who take "control" to mean ONLY stabilization of variables subject to disturbances. Then they see the sort of claims you describe as nonsensical -- see my reply to Bill Powers below.

>(Incidentally, I'm still very uncomfortable with the conflation of
>"behaviorist" and "S-R viewpoint."

I've railed about this for years. May your discomfort be heeded more than mine! (Actually, I think there has been fair progress among PCTers over the past couple of years in developing critiques of Skinnerians.)

>As I like to use the word, any psychologist who is not a behaviourist should
>be considered a mystic. PCTers are not mystics. They observe what happens
in
>the world and model it using plausible, possible mechanisms that are as
>similar as possible to what is known from real observations of real people
and
>physically described machines. PCT is a thoroughgoing behavioural theory, in
>my view. And I think you know by now that I do have a slight inkling of what
>PCT is all about.)

So you would count virtually all of contemporary cognitive psychology as "behavioristic"? Skinner wouldn't have, because he didn't think modeling the CNS innards is worthwhile. Did that make him a SUBSPECIES of behaviorist? If so, that SUBSPECIES is (at least one of) the foes of PCT. I myself prefer to think of PCT (and especially HPCT) as "mechanistic," which (to my notions) describes both behavioristic and cognitivistic methodologies -- and PCT in a sense combines, and so extends, those methodologies.

>Bill Powers (920826.2000)

>There's one problem with the "self-consistent" behaviorist scheme that
>you didn't mention: the reference signal. When a behaviorist trains a
>rat to do something, what the rat is to do is already known to the
>experimenter. As I've said before, a behaviorist who wants to
>demonstrate the reality of stimulus control of behavior must state
>BEFORE the demonstration what behavior is going to be brought about.
>In other words, the experimenter must describe what will be perceived
>when the process has finished.

I meant to include that problem in (1) not having predictive models.

>When Skinner spoke about an experimenter controlling an animal's
>behavior, his language always contained some way of saying that the
>experimenter desired, specified, or intended that a certain behavior
>come into being. In talking about the environment controlling
>behavior, of course, he did not use such terms.

But Skinner was careful to say that ANY implication of "purpose" in his writings is to be construed as historically determined by the organism's history. That holds for experimenters as well as for subjects. The claim is that "purpose" always can be unpacked to "the organism's past resulted in the organism now doing..."

>Only the experimenter was allowed to behave in a purposive way, knowing what
>was to happen and doing whatever was required to make it happen.

No, Skinner (when being careful and not playing to the lay crowd) would have been happy to dispense with "purpose" entirely, and say that ALL organisms behave for the same (historically determined) kinds of reasons -- and even say that purpose (no quotes) is an ILLUSION. (Ditto knowing and doing what is required and intending and wanting.)

>Therefore he attributed true control only to behavioristic experimenters, and
>not to subjects or to environments.

No, he NEVER attributed "true [PCT!!!] control" to ANY organisms -- ONLY the ability to alter the environments of other organisms. And the playing out of that ability in any given case is set by the organism's history.

>He [Skinner] was doing as he often did: he used a common word that to most
>people implies intention or purpose, but with the understanding that only his
>meaning was allowed. Something like the way certain politicians are now
>talking about "family values."

Yes, he did do that. And you, as a PCTer, use the common word control with the understanding that only YOUR meaning is allowed -- and don't tell me that it is EXACTLY the same meaning as used by people-in-the-street! (It is closer than Skinner's meaning, but not identical, I say.) So, never the twain shall meet.

>So I dispute use of the term "self-consistent" in reference to
>behavioristic notions of control.

As I show above, Skinner IS self-consistent. At the expense of restricting his meaning of "control" to something not acceptable to YOU. I think he (maybe) could have seen that PCT is self-consistent at the expense of restricting ITS meaning of "control" to its technical definition. The (contrasting) positions are analogous. But PCTers must be as careful in their own way as Skinner had to be careful in his -- sloppy talk about "autonomy," for instance, must be grounded fully by going back to the definition of control. Remember our discussion awhile back about the lack of traditional "free will" in PCT systems?

Greg

Date: Thu Aug 27, 1992 8:34 am PST
Subject: Demonstrations

I would like to get dem1a, dem2a, the blindmen.doc and all three marken files outlined in the introductory letter. Unfortunately I have not learned much Unix and don`t know how (or if it is even possible from here) to get (ftp) these files. I am working on a sun in globalview but I can get to unix if I have to. I am planning to read and play arround with the files before school starts (at home in dos). As my employment ends on friday, that being tomorrow, (as with my access to this terminal) I would really appreaciate a quick response.

I have also heard about other demo`s (pointing finger, five digits, mind reader, etc) if these are available I would like to try them also. Please, only control stuff for now. I would like to learn to think this way a teeny weeny bit better before taking on all commers.

-John
-Student of electron behavior in materials

John van Loon: John_Van_Loon.XRCC%xerox.com@uunet.ca

Date: Thu Aug 27, 1992 9:58 am PST
Subject: SR; Interactionism;Agre's Cooking program

[From Bill Powers (920827.0800)]

Mary tells me that Living Control Systems II is in "Books received" this week in Science.

Martin Taylor (920826) --

RE: Control of other control systems

You say
>What is the problem with controlled agents themselves controlling others?
>Let us take the S-R position that some pattern in the environment determines
>the behaviour of an agent (I assume this is what "control" means in that
>context). Then why should the actions of that agent not create a pattern
>in the environment that determines the behaviour of another (sub)agent?"

>It seems to me analogous to the situation of a high-level ECS controlling
>its percept by way of (unknown to it, of course) setting references that
>lower ECSs use in controlling their percepts.

The most important disanalogy is that one organism has no way of directly setting reference signals for another. At best, one organism can disturb another's controlled variables, and thus create some sort of shift in the lowest levels of perception. As a result, there may be deviations of higher-level perceptions from their reference signals, countering which require changes in lower-level reference signals. The nature of those changes depends on ALL the control systems in the organism, not just the system disturbed, and is unpredictable without complete knowledge of the organism's hierarchy of control. As your own theory of Layered Protocols makes so clear, all communication between independent organisms must take place through the lowest levels of the system. That is the only contact between the higher levels.

Inside a single organism, of course, higher systems use lower ones quite freely, just by manipulating their reference signals through the natural connection from one level to another. Lower levels have no way of adjusting their own reference signals. But no such connection exists between people; all people adjust their own reference signals at all the levels where other people can do the same thing.

Also, I should ask you why we should take the S-R position that patterns in the environment determine the behavior of an agent. To say this implies that given only the pattern, we can predict the behavior of the agent. I agree that this expectation follows from the S-R premises. But I don't agree that the expectation has ever, in fact, been fulfilled. All that patterns in the environment can do is change perceptions. What is done about such changes depends on the organism's goals, not on the pattern.

If the relations between external patterns and behavior were deterministic, we would expect a clear relation of dependency to be observed. Instead, we see that there is always a large scatter, of the same magnitude as or larger than the magnitude of the average apparent effect. This is hardly a "deterministic" relationship.

>... whether you buy this analogy or not, I can't see the paradox when
>the controlled controller is seen strictly from an S-R viewpoint.

The paradox can't be seen from a strict S-R viewpoint. I'm very puzzled by these words of yours. I certainly agree that you "do have a slight inkling of what PCT is all about." But the impression you're creating here is that there is something about the S-R explanation of behavior that you still find valuable. Is that true? Or am I once again misinterpreting?

Penni Sibun (920826) --

I'm being cautious here to avoid more podiphagia.

You say
>the interactionist argument is in general deemphasizes individuals and
>emphasises their interactions.

Does this mean that interactionists believe that interactions can be EXPLAINED without reference to the properties of individuals, or that they can be DESCRIBED without such reference? Does the interactionist viewpoint differ from, say, that of conventional sociology, in which it is held that there are laws of social systems that exist apart from, and indeed unconnected to, laws of individual behavior?

Avery Andrews (920826) --

>Bill's arm demo (the version I have at least) has to assume that the
>visual system can track the fingertip & the target, whereas a major
>aspect of Sonja is a visual system that can to some extent actually
>do this (though he ducks problems of early vision).

In the Arm Demo, I assume that the visual system can identify the target and the fingertip as discrete objects in the retinal image. I assume that it can generate signals that depend on the retinal separation of these objects in the x and y directions. The retinal separation is given in units of angle as seen from the position of each eye's lens, which is proportional to the spatial separation of points in the retinal image.

Can you explain how Sonja's visual system identifies the various retinal images, and how the relative retinal distances between image points are converted into signals?

-- I have a copy of a '92 paper for AAAI (whatever that is) by Agre and Horswill called "Cultural support for improvisation." In it he gives the results of a run of a program called "Toast," which makes breakfast. Evidently the output is a series of statements indicating what is being done by the program at each step.

```
0      break-egg egg-11 bowl-4
       [making omelette]

1      add-egg egg-10 omelette-batter-0
2      add-egg egg-9 omelette-batter-0
3      beat omeletter-batter-0 whisk
4      move pan-4 burner-4
5      move butter-pat-15 pan-4
6      melt burner-4 pan-4 buttter-pat-15
7      move slice-23 toaster\
       [Waiting for butter so making toast]

8      start toaster slice-23
9      move knife-4 plate-1
10     *** done with goal (knife clean plate-1)
```

11 move fork-4 plate-1
12 *** done with goal (fork clean plate-1)

13 move butter-pat-14 knife-3
[back to toast]

14 butter slice-23 knife-3 butter-pat-14
15 pour-omelette-batter omelette-batter-0
[butter melted so back to omelette]

16 move slice-23 plate-1
[setting table]

At the end is:

> Figure 2. Sample run of the breakfast program. The agent was given the goals of making an omelette, two pancakes, a slice of toast, and setting the table, then cleaning up. Our comments appear in square brackets.

Then, in the section called "Control experiments":

> To insure that the breakfast maker problem was not entirely trivial, we implemented a STRIPS formalization of a subset of the domain and tested it using the SNLP non-linear planner [ref] on the omelette making subproblem. A solution to this involves breaking and beating three eggs, melting butter, initial cooking, then folding and browning each side. The planner could solve the problem when led through the correct plan a few steps at a time, but could not solve the complete subproblem from start to finish given the resources which we had available to us.(The planner ran for 6 hours on a Symbolics XL1200 before exhausting its paging disk). The planner was able to make toast and set the table however.

If this is what is meant by "too much machinery," I get the point. 6 hours!

Before this section, the authors discuss the philosophy of their approach, which depends on using the properties of objects in the environment to replace a great deal of computation. Here's my interpretation of what they say; make allowances for the fact that I'm trying to swim through a sea of unfamiliar terminology. Gratuitous comments by me are in square brackets.

Objects in the world are characterized by the roles they play in customary routine activities. [The playing of roles is attributed to the objects. A lot of the language that follows could be cleaned up by inserting "are perceived to be" or some such proviso at frequent intervals].

A given object type is characterized by a state graph: nodes are states and arcs are labeled as operations that change the state from one node to another. [These are apparently thought of as properties of the objects, but operations are defined in terms of their outcomes, not in terms of the actions that bring them about. In HPCT terms, the operations could be re-phrased as the setting of goals specifying a new state which lower systems then bring about, thus moving the object state to the new node].

An operation can be a label that applies to several arcs leaving a given node, thus being "nondeterministic." [This baffles me. Do they mean "stochastic?" Do they mean that under various circumstances, a given act can have different outcomes? We begin to see the penalties of failing to distinguish between the actions of the system on an environment and the response of objects in that environment to the actions. Properties of the behaving system are being merged with properties of the environment.]

A object instance has an associated object type.

A world state is a set of instances and a mapping from each of these instances to a state in its type's graph. A history follows a set of instances through a series of changes. [This would seem to amount to a model of the environment's behavior]

An action type is a sequence of pairs made of an object type and an operation from that type's state graph.[So an action type is a set sequence of operations that produces a set sequence of states. The assumption that there are no disturbances is clear. A given type of action will always produce a given result. The need for control theory is evident.]

A history is thus determined by an initial world state and a sequence of action instances.

A task is a set of object types, an initial world state, and a set of goals, each of which consists of an object type and a state from that type's graph. [So reference levels are picked from states known to be possible for that type of object -- see the role of memory in the HPCT model.]

A goal for an object is satisfied if the world state contains some instance of the object type that is in the indicated state.(So one cannot assert that two instances of a given type must be in a given state). [I don't see the advantage of this rule. It says, apparently, that the goal of breaking a given egg is satisfied if the egg being operated on remains intact, but another egg breaks. But I must remember that the environment in this system can't do anything spontaneously -- there are no disturbances. So the only egg that can break is the one being operated on by the break operation and there's no need to keep track of which egg is broken.]

The authors say of this scheme as described so far, "... it allows us to define one important source of complexity-ameliorating constrain in a natural way, as follows."

What follows is a series of categories. First, a class of tasks called cooking tasks, containing two object types: tools and materials.

A tool is an object that possesses at least one free state (meaning that this state can be reached from any other state using a single operation). Each tool has a normal state, such as "clean and dry."

A tool action involves one or more tools; a normal tool action requires the tool to be in its normal state at the beginning.

A material has a raw state, from which all its other states are reached by operations in the class normal tool actions. [This reminds me of writing a program in which all the variables are arrays of pointers to structure entities, in anticipation of later expansion to a general case. In my experience, the programs always change before the general case is reached, so the program becomes excessively ponderous because of including generalities that turn out never to be used.] Other operations on materials can also occur.

There is also a third object class called containers.

Well, we can see roughly the universe in which this approach takes place. There are parts of it that suggest how level-9 (program) behaviors might be organized, with sequence and category being included, but all lower-level control processes being assumed to take place successfully. That assumption is OK with me. What's missing, of course, is any explicit recognition of the goal-directed tasks as being control processes. The ways of categorizing the elements of these tasks are arbitrary; more useful categories, I think, would result from including one lower level of perception, relationships. Even just "in" and "on" would be enough to eliminate the category "container." To go on:

"We assume that the agent can readily detect the states of all visible objects. The agent achieves its goals by using normal tool actions to push materials through their state graphs along certain customary trajectories; and it uses focused actions to get tools into normal states. The algorithm thus depends on a culturally defined abstraction that minimizes interactions among the actions in cooking tasks."

Summing up, the authors give this description of the algorithm:

While there are unsatisfied goals,

Choose aribtrarily an unsatisfied goal.

Let G be the specified goal state.

Find an instance M of the specified material type.

Determine its current state S.

Let A be the entry for the S and G in M's action table. Choose tool instances to fill each of A's participant slots, preferring instances which are already in their normal state.

If some chosen tool T is not in its normal state, then perform the action specified in T's action table for its current state. Otherwise perform A using M and the chosen tools.

This last summary is ever so much simpler than all the verbiage that led up to it, because the attempt to present a general case has been reduced to a specific case. We can see how to relate it to control theory rather directly, because the algorithm is basically a control process. A goal is given, compared with a current state of affairs, and steps are taken to reduce any difference between the goal-state and the perceived state to zero.

One change I would suggest is the explicit addition of an error graph for each operation. The main weakeness of this algorithm as it stands is that every operation must produce a predictable state. By using an error graph this limitation could be removed.

There are many states in which an object can be, only some of which are needed at a given moment, the rest being departures from that state. For each perceived state of an object and each goal state of the same object, there is an error-state. The action that is required to decrease that error depends on how much error there is, and what its direction is. If, for example, a sub-goal is "break three eggs into the pan", and there are no eggs left on the kitchen counter after two have been broken, the lower-level action that brought the eggs to the counter in the first place should be used again, with an argument of 1: "get 1 egg and put it on the counter." This reduces the error in the number of eggs to zero and allows the higher-level process to complete the correction of its error.

The error graph could be very complex unless we reduce each perceived variable to a single dimension (like number-of-eggs-on-the-counter), multiplying the number of systems operating in parallel as required. When the perceptual variable can vary only in one dimension, the error graph becomes one-dimensional: for a particular perception and a particular goal, pick a subgoal that will increase or decrease the difference as required. It is not required that the action eliminate the error; only that it make the error smaller. If the error gets smaller on every iteration, it will eventually go to zero. In many cases it can go to zero in one or two iterations. This allows for disturbances to occur, both from independent events and from the actions of other control systems in the same agent operating at the same time. A disturbance will simply move the error graph to a new node, which will indicate the sub-goal adjustment needed to reduce the error from that starting point. So if the agent accidentally drops an egg, it will reactivate the lower-level system that got the other eggs from the refrigerator and complete breaking the required number of eggs into the pan.

This is saying no more than that this algorithm, since it already resembles a control process, might as well be carried the rest of the way.

Best to all, Bill P.

Date: Thu Aug 27, 1992 9:59 am PST
Subject: Misc Replies

[From Rick Marken (920827.1000)

Martin Taylor (920826 14:45)

I said:

>OK. There's the trivail truth.

Martin said:

>Even when you acknowledge the
>fundamental truth, its consequences and power are far from self-evident.

For sure. But Martin, you never answered my question; namely, what is the deeper (non-trivial) truth revealed by dynamic analysis?? I really don't understand what you mean. I know that control systems are dynamic and that these dynamics are important to consider when analyzing and building control systems. But can you give me a concrete example of what you mean by a non-trivial truth. I kinda think it's non-trivial to realize that

1) behavior is just an observable side-effect of an organism's efforts control various aspects of its own sensory experience relative to internally specified reference levels and that

2) this means that the only way to understand behavior (non-superficially and non-statistically) is to try to learn what sensory variables the organism is controlling.

These two trivial little truths imply a hell of a big research program. What truths did you have in mind? Concrete examples would help.

Martin says:

>I'm afraid I don't see the paradox. What is the problem with controlled
>agents themselves controlling others?

Bill answered this for me. I will just add one little thing. There is, indeed, really no paradox involved, though it sounds like there is one when it is stated as I stated it above (how do controlled agents control?). The fact of the matter is that agents that control cannot be controlled -- and efforts to control controlling agents are simply examples of conflict between control systems.

Avery says:

>Moving on to a more specific point, Bill's arm demo (the version I
>have at least) has to assume that the visual system can track the
>fingertip & the target, whereas a major aspect of Sonja is an
>visual system that can to some extent actually do this (though he ducks
>problems of early vision).

I don't understand this. The arm demo's visual system does track the fingertip and target, doesn't it? What does Sonja do about target tracking?

Bill Powers says:

>Avery, you are becoming the voice of reason on this net. Keep it up.
>You will get us back on track.

Not if I can help it! (Just some unreasonable sarcasm). You are a remarkably temperate and reasonable control system, Avery, just like your Aunt Mary.

Bill to Martin RE: degrees of freedom.

>I think we will all continue to be frustrated on this subject until we
>come up with experiments that demonstrate this "absolute need."

Ah, the same unreasonable suggestion that I made.

Bill to Ed:

>I put it to you that PCT and HPCT transcend political points of view, as
>well as many others. Except for Rick. Yukayukayuk.

It's true, folks. I AM LEVEL 11.

Martin Taylor (920827 0020) answers Bill's suggestion that experiments are needed to demonstrate the need for revision of the model based on degrees of freedom considerations:

>I doubt you could ever find a phenomenon that can be described by only
>one organization.

True. But a collection of observations can usually only be explained just one of the models. All I'm asking for is some data that would motivate the need to revise the existing model (which handles many simple cases of control). I would suggest a very simple experiment; just so I could get the idea of what you are talking about. tI scould probably just be a simple extension of a multidimensional tracking task (which we know the model can handle). Just create a simple change in the experiment that produces consistent results that pose a problem for the existing model.

>There really isn't any experimental demonstration that I can think of,
because

If you really can't do an experiment that tests your proposed revision of the model then I can't see that it's such an important revision; since the model accounts for the phenomena we've observed so far I can't see revising the model to help it explain a phenomenon that we will never be able to observe.

Hasta luego Rick

Date: Thu Aug 27, 1992 10:23 am PST
Subject: Perceiving as a PCT theorist

What is the relationship between the charts for each teacher/councelor/participant called for here & the kinds of charting called for in Demming Mangement TQM sorts of things?

Date: Thu Aug 27, 1992 10:32 am PST
Subject: Perceiving as a PCT theorist

from Ed Ford (920827:1015)

I'm leaving early tomorrow morning for my son's in California and I'll be back Sunday afternoon. I'm still overwhelmed from last week's absence and will address Mark's thoughts and David's request early next week. Below, David, is part of what I was doing.

>(Powers 920826) No, Ed, you are not a Republican. You are a a control theorist.

And proud of it.

Concerning control theory and modeling......

I am not a scientist and I don't build models in computers. However, as Bill mentioned, I am a control theorist. When I work in any setting, be it counseling, a residential treatment center, with organizations, individuals, especially schools, when I try to address the problems presented I do so as a control theorist. That means I perceive a whole bunch of living control systems (or one individual control system) trying to get along with other living control systems, whether it be a home, school, or whatever. When a friend meets and talks with my wife, Hester, he perceives her as a woman, mother of eight, and my wife. When I perceive Hester, it is quite different because there are different factors that go into making her a perception which are used by me to create the concept Hester in my mind. I think the same is true no matter how we address any social problem. Obviously, to think as a control theorist you have to understand the basics of the theory, not to the extent that Bill and most other members of the CSG do. But it is critical to understand the basics.

I found this control theory way of perceiving things very helpful last week when I was working with teachers and parents in the Yakima School District. As Bill mentioned elsewhere in his post, the Outcomes driven model is nothing other than controling input or, controling perception. All the individuals set their own goals and the system is driven by the outcomes (perceptual variable). Jack Champlin, who conceived the model, obviously didn't make the connection between that and control theory because he didn't understand PCT (although he is getting there). As a control theorist, once I understood what he was saying and doing, I was able to help him take his model further.

For example, one of the problems he has been thinking of is how do you statistically keep the system honest. I suggested that you have to set very specific reference signals or goals, then you set up a chart and measure over time how you presently perceive how you are doing with relation to what your specific goal is. Everyone in the system, from the school board, superintendent, on down through the school administrators and teachers to the students; everyone should have their own individual chart for measuring how they are doing. That is the only way I can figure how all these individual control systems, all with their own ways of perceiving things, with their own sets of priorities, standards, and values, can possibly work together. They have to all agree to a common over all goal, then define their own individual goals that, if achieved, would move the organization to achieving it's goal. It keeps the whole system honest. This has another benefit, by making use of a visual chart that reflects the historical progress, the actions are being driven more precisely, and the system can deal with it's goals more easily and, more importantly, will continually know how it is doing. If you are driving a car and the speed sign says "drive safely", what speed do you go. Or, if the sign says 55 MPH and your speedometer isn't working, you'll still have problems because you'll know what you want but you will have difficulty comparing it to how you are doing. Thus the need for specifics. Thinking as a control theorist gives this advantage.

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

Date: Thu Aug 27, 1992 10:39 am PST
Subject: Re: SR; Interactionism; Agre's Cooking program

[Martin Taylor 920827 13:30] (Bill Powers 920827.0800)

>You [MMT say
>>What is the problem with controlled agents themselves controlling
>>others? Let us take the S-R position that some pattern in the
>>environment determines the behaviour of an agent (I assume this is
>what >"control" means in that context). Then why should the actions
>of that >agent not create a pattern in the environment that determines
>the >behaviour of another (sub)agent?"
>
>>It seems to me analogous to the situation of a high-level ECS
>>controlling its percept by way of (unknown to it, of course) setting
>>references that lower ECSs use in controlling their percepts.
>
>The most important disanalogy is that one organism has no way of
>directly setting reference signals for another.

Yes, but remember that reference signals have no place in an S-R view of the world. Rick's original question posited an intrinsic paradox in the S-R view, that should, in itself, be enough to get S-R people to see the error of their ways.

>>... whether you buy this analogy or not, I can't see the paradox when
>>the controlled controller is seen strictly from an S-R viewpoint.
>
>The paradox can't be seen from a strict S-R viewpoint. I'm very
>puzzled by these words of yours. I certainly agree that you "do have a
>slight inkling of what PCT is all about." But the impression you're
>creating here is that there is something about the S-R explanation of
>behavior that you still find valuable. Is that true? Or am I once
>again misinterpreting?

"Wud some Guid the giftie gie us,
Tae see oursens as ithers see us."
(Rabbie Burns --more or less)

All I am saying is that when looking for paradox, one should not take two simultaneous viewpoints. I do not in any way accept the S-R position. But to someone who does not see its intrinsic nonsensicality, there is no problem with the controlled controller at all. There's no paradox. Neither, in PCT terms is there any control.

So don't worry about whether I support the S-R position. It is intrinsically unsupportable, and needs no spurious paradoxes to show its problems.

Martin

Date: Thu Aug 27, 1992 11:05 am PST
Subject: Re: Misc Replies

[Martin Taylor 920827 14:00] (Rick Marken 920827 10:00)

I'm lost. I don't understand what it is that Rick doesn't understand so I can't answer his questions. Rick seems to be saying that there are no more sensory degrees of freedom than motor degrees of freedom, and that all ECSs can always be controlling. That, at least, is what I think is the case in the models I have seen simulated or discussed. But it can't be true, and I'm sure Rick knows it can't be true, any more than it can be true that we behave in a strict S-R manner. So, since I believe Rick to be reasonable, there must be some intent behind his questions that totally escapes me.

=============quotes follow.

All I'm asking for is some data that would motivate the need to revise the existing model (which handles many simple cases of control). I would suggest a very simple experiment; just so I could get the idea of what you are talking about. tI scould probably just be a simple extension of a multidimensional tracking task (which we know the model can handle). Just create a simple change in the experiment that produces consistent results that pose a problem for the existing model.

>There really isn't any experimental demonstration that I can think of,
>because

If you really can't do an experiment that tests your proposed revision of the model then I can't see that it's such an important revision; since the model accounts for the phenomena we've observed so far I can't see revising the model to help it explain a phenomenon that we will never be able to observe.

============a sort of response? follows:

Has the model accounted for all the phenomena of human life observed so far? I thought that I mentioned a couple of phenomena that seem not to be covered by the N-in N-out model: the flick of the eyes to an unexpected movement, the turn of the head toward an unexpected noise. These seem to involve a K-in N-out model, where K>>N. This seems crystal clear, and I don't suppose that's where the problem lies.

All I am doing is simply pointing out a fact that must be true in the real world (as opposed to the world of constructed models acting in trivial environments). It's as much a fact as that disturbances sometimes affect CEVs that correspond to percepts. The latter fact leads to PCT in general, the df discrepancy leads to what I have called "alerting systems" within PCT.

Perhaps there's a misunderstanding about the difference between being able to observe a phenomenon and being able to do experiments on it. Rick, is that where the problem lies? Astronomers can observe but not experiment. They can model their phenomena using the results of experiments done by physicists and chemists. My claim was not that there are no observable phenomena attributable to alerting systems, but that we couldn't undo evolution, and wouldn't ever find an active living system without them; and if we constructed a simulation, its construction would be whatever the designer decided. If it failed to respond to an alerting condition, the reason would be in the design.

And I did ask if anyone else could think of an experimental situation that might tease out the mechanisms of alerting (as opposed to its existence). Maybe there could be experiments in the area of high workload stress, but they would probably be very complex, and I can't see how they would work.

Finally:
>>I doubt you could ever find a phenomenon that can be described by only
>>one organization.
>
>True. But a collection of observations can usually only be explained just one
>of the models.

When you are comparing the relative merits of a fixed set of models, any observations can change their relative likelihoods. But there is always an infinity of other models standing in the wings awaiting their turn to be tested.

I have a fundamental proposition about science that I have seen no reason to change since before my grad school days:

"Whatever your theory, it is either imprecise or wrong."

Martin

Date: Thu Aug 27, 1992 12:13 pm PST
Subject: The paradox of control

[From Rick Marken (920827.1230)]

Martin Taylor (920827 13:30) says;

>All I am saying is that when looking for paradox, one should not take two
>simultaneous viewpoints.

But that is precisely what the Behaviorists (a la Skinner) are doing!! In order to avoid the paradox, they use "control" in the PCT sense when they claim to be able to control behavior, and they use control in the "causal" sense when when they say that behavior is controlled by the environment.

>So don't worry about whether I support the S-R position. It is intrinsically
>unsupportable, and needs no spurious paradoxes to show its problems.

Let me try to explain what the paradox is and why I think it is not only not spurious but genuinly important (worth meditating on a bit).

The paradox is like that of the man who claimed to be from Crete, where all men are liars.

Skinner is saying that people are controlled and that he (a person) can control them.

If he means control in the PCT sense then this is paradoxical because it implies that he can "control a person" (bringthe person's behavior to an intended state) even though his (Skinner's) own behavior is controlled and, hence, being brought to an intended state that may be different than "control that person". If Skinner can control (meaning that he can select a particular behavior for a person to exhibit) then he cannot also be controlled (which would open the possibility that he would have no choice but to do some- thing other than select THAT particular behavior for the other person to do). So, for Skinner's claim to be true (and non-paradoxical) Skinner would have to be the only person who is NOT controlled (in the PCT sense), in which case there is no possibility that that which is controlling Skinner could get Skinner to do anything other than control the other person. The same applies to the Cretin; for his statement to be true, he would have to be the only non-lying Cretin.

If Skinner means "control" in the causal sense then the paradox does seem to evaporate. But then it seems a bit ridiculous for Skinner to encourage us to control other people -- since this will happen anyway and the results of our influence will be what they will be -- we have no PCT control over the outcome. But Skinner is not using control in this way (obviously) when he claims to be able to control people (as Bill explained) and when he recommends that we control people in order to produce a particular result; viz. a good society full of well behaved people.

So I think the paradox of control is very informative. It helps one get their arms around what it really means to control (which is the PCT meaning -- ie. produce an intended outcome despite the existence of influences on that outcome that would tend to make it different than intended). Control (as Skinner understood about his own efforts at control) is made of sterner stuff than are causes and influences. Control is VERY purposeful -- in fact, it is purpose.

The paradox of control reveals the fundemental problem with any scheme for improving the world by controlling people. If you attribute to all people the same capabilities that you assume for yourself (the ability to control people) then you are denying yourself the very ability you claim. For if you can control, then you cannot be controlled -- and neither, then, can those who would be the objects of your control.

Best regards Rick

Date: Thu Aug 27, 1992 12:51 pm PST
Subject: Re: SR; Interactionism; Agre's Cooking program

(ps 928027.1300)

this is just a placeholder that i will comment more on bill's msg. when i can (this list goes by so fast!), but basically i think his analysis of toast is excellent, and shows real promise of a rapproachment bet. pct and interactionism.

more later! --penni

Date: Thu Aug 27, 1992 1:28 pm PST
Subject: Re: The paradox of control

[Martin Taylor 920827 16:15] (Rick Marken 920827.1230)

I think I agree with the paradox of control when seen from a PCT viewpoint, with the caveat that full control implies infinite gain, and any departure from that leaves a little hole in the argument. It's a bit like the liar's paradox that you used as an analogy. That, too, doesn't work when there are degrees of truth-telling rather than pure lies and pure truth. But let that be.

You presented the paradox as a way to get non-PCT believers in S-R to see that their position was untenable and that they should convert to the true way. If the paradox evaporates when seen from the S-R position, then it does not provide a very powerful means for inducing that conversion. That's what I tried to illustrate, in the message that distressed Bill.

When you try to affect someone's actions, it helps to try to see things from their viewpoint. Of course it isn't possible to do so with any precision, but one can often get close. Good advertisers and propagandists get quite close. People who start with the position that they are correct, and that others should see it from the "correct" point of view, do not get close. They may well be correct, but they won't have much influence.

Martin

Date: Thu Aug 27, 1992 3:21 pm PST
Subject: Straw Skinner

From Greg Williams (920827 - 2) >Rick Marken (920827.1230)

>In order to avoid the paradox, they [Skinnerians] use "control" in the PCT
>sense when they claim to be able to control behavior, and they use control in
>the "causal" sense when when they say that behavior is controlled by the
>environment.

If my previous post of today fails to convince you that this is incorrect, perhaps you had better consult Skinner's writings, e.g., SCIENCE AND HUMAN BEHAVIOR, p. 228: "When a man controls himself, chooses a course of action, thinks out the solution to a problem, or strives toward an increase in self-knowledge, he is BEHAVING. He controls himself precisely as he would control the behavior of anyone else -- through manipulation of variables of which behavior is a function." Read the sequel on "self-control" at your leisure. Suffice it to say, once again, that Skinner did NOT "use 'control' in the PCT sense" AT ALL. Q.E.D. (or R.I.P., or maybe not so peacefully, since the PCT bulldozer is headed Skinner's way!)

>Skinner is saying that people are controlled and that he (a person) can
>control them.

...

>If Skinner means "control" in the causal sense then the paradox does seem
>to evaporate. But then it seems a bit ridiculous for Skinner to encourage
>us to control other people -- since this will happen anyway and the results
>of our influence will be what they will be -- we have no PCT control over
>the outcome. But Skinner is not using control in this way (obviously)
>when he claims to be able to control people (as Bill explained) and when
>he recommends that we control people in order to produce a particular
>result; viz. a good society full of well behaved people.

Skinner said that ALL interacting organisms control (that is, alter the environments of, and so contribute to altered histories of, and hence altered behaviors of) each other to various extents. He also claimed that one's efficiency of control (in the popular sense of making other organisms behave in certain ways) can be enhanced by certain techniques which he (Skinner) was promoting. Suppose, Skinner would say, you behave as if you "want" another organism to behave in a certain way; that behavior is the result of your history. If your history were different, you might not behave as if you had that "want." Now suppose you DO behave that way. Then, for your control (Skinner's meaning) to result in efficacious control (popular sense) might require special techniques (which Skinner claimed to have at hand). In fact, his techniques work best when the controllee isn't in the same league smarts-wise as the controller. In old age, Skinner complained about verbal behavior among equals as more or less doing in his hopes for easy control of humans based solely on his ideas developed via experiments with rats and pigeons; so your comments at the end of your post regarding the DIFFICULTY of Skinner's (non-PCT) control programme are well taken.

I think you need to read what the opposition (loyal or not) says before putting words into their mouths. Your turn to quote scripture.

Greg

Date: Thu Aug 27, 1992 3:25 pm PST
Subject: Re: The paradox of control

[From Rick Marken (920827.1530)]

Martin Taylor (920827 16:15)

>You presented the paradox as a way to get non-PCT believers in S-R to see
>that their position was untenable and that they should convert to the true
>way. If the paradox evaporates when seen from the S-R position, then it
>does not provide a very powerful means for inducing that conversion.

I'm not that interested in the paradox as a way of getting S-R theorists to switch to control theory. I am much more interested in the paradox as a way of illustrating the self-contradictory nature of the assumption that human problems can be solved by control of other people. This solution to problems is not just a penchant of SR theorists or reinforcement theorists. It is the the approach that is advocated (and used) by many people, every day, as a way to make the world a better place. I bet that most people think about control just the way Skinner did.

They believe that people can be controlled and that they themselves can exert that control if they just try hard enough. Parents believe it; teachers believe it; civic leaders believe it; employers believe it; managers believe it; children believe it; even I apparenly believe it sometimes (as some of my efforts to brow beat people into accepting PCT might suggest -- thanks Chuck). The paradox of control is for people who believe that they can make things better (ie. the way they want it) by controlling other people. It points to why it doesn't work even though it seems like it should. It suggests how important it is for psychologists to start trying to understand how people control, and to stop worrying about how to control people.

Best regards Rick

Date: Thu Aug 27, 1992 4:02 pm PST

[Avery Andrews 920828.0923] (Bill Powers 920827.0800)

>You say
>>the interactionist argument is in general deemphasizes individuals
>and emphasises their interactions.
>
>Does this mean that interactionists believe that interactions can be
>EXPLAINED without reference to the properties of individuals, or that
>they can be DESCRIBED without such reference?

Neither, based on what I've seen. In the present intellectual climate out there in ScienceLand, it is just a matter of whether you bother to think about interactions at all. E.g. all PCT-ers are already Interactionists.

Avery.Andrews@anu.edu.au

Date: Thu Aug 27, 1992 4:33 pm PST
Subject: sonja's visual system

[Avery Andrews 920828.1013] (Rick Marekn 920827)

The short answer to what's so great about Sonja's visual system is that she has a simulated visual field (as would be provided by current ideas about low level vision), in which objects are located, tracked, etc. by means that are (supposedly) plausible from a neurological point of view.

Avery

Date: Thu Aug 27, 1992 6:28 pm PST
Subject: arm demo questions

[Avery Andrews 920828.1209]

As a sort of comprehension check on my understanding of the arm demo (old version), is there a reason why it was done the way it was done, rather than the way I would have done it, having position error-signals set reference-levels for joint-angle velocities (I don't claim any advantages for my way--it's just what would occur to me first). Is there any known relevant empirical evidence?

Also, I'm unclear on why there needs to be a computation for a kinaesthetic vertical angle (A + B/2, where A is the vertical shoulder angle and B is the elbow angle). Why couldn't the vertical shoulder angle reference level be derived directly from the visual system?

Avery.Andrews@anu.edu.au

Date: Fri Aug 28, 1992 4:32 am PST
Subject: No Paradox of Influence

From Greg Williams (920828) Rick Marken (920827.1530)

>I'm not that interested in the paradox as a way of getting S-R theorists
>to switch to control theory. I am much more interested in the paradox
>as a way of illustrating the self-contradictory nature of the assumption
>that human problems can be solved by control of other people. This
>solution to problems is not just a penchant of SR theorists or
>reinforcement theorists. It is the the approach that is advocated (and
>used) by many people, every day, as a way to make the world a better
>place. I bet that most people think about control just the way Skinner did.
>They believe that people can be controlled and that they themselves
>can exert that control if they just try hard enough.

Once again, Skinner's "control" = PCT's "influence" -- consistently. If "most people think about control just the way Skinner did," then they believe that people can be influenced and that they themselves can influence people to have particular outputs. And that belief IS in fact true, isn't it, at least in a circumscribed way? Advertising, propaganda, education, those beautifully contrived con games, "psychic" demonstrations, politics, family interactions all show it to be true, I think. Do you disagree? Are we incapable of influencing each other at all? As I've noted before, PCT tells us that the more we know about others' control structures, the better we are able to manipulate them to do what we want -- just set up situations so they will be satisfying their reference signals by doing what we want (and they needn't even be aware of being manipulated). So PCT-inspired manipulation could be quite insidious! (It could even get worse. One's current reference signals could result largely from past manipulations. "My country, regardless," etc.)

>Parents believe it; teachers believe it; civic leaders believe it; employers
>believe it; managers believe it; children believe it; even I apparently
>believe it sometimes (as some of my efforts to brow beat people into
>accepting PCT might suggest -- thanks Chuck). The paradox of control is for
>people who believe that they can make things better (ie. the way they want
>it) by controlling other people.

If "control" means "influence" -- and I think it does for virtually everybody, there is no paradox; people (even you) believe it because it works. People CAN (to an extent) make things the way they want by influencing others.

>It points to why it doesn't work even though it seems like it should.

It DOES work in a circumscribed way. And PCT shows why. PCT explains why influence which opposes the influencee's wishes is resisted. PCT also explains why influence which is consonant with the influencee's wishes is NOT resisted, in fact is embraced. In short, PCT explains why influencers need to be sophisticated about their techniques to be successful (in particular, it says that they should learn about the control structures of the influencee -- and advertisers and con men understand this to the nth).

>It suggests how important it is for psychologists to start trying to
>understand how people control, and to stop worrying about how to control
>people.

Understanding what a particular individual is trying to do can be the key to successful manipulation, but I don't think that is how you meant the statement. And given that there is no paradox of influence, there is no reason for psychologists to stop worrying about how to influence people.

Greg

Date: Fri Aug 28, 1992 5:39 am PST
Subject: Re: steering

Forgive -- if need be -- this humorous interlude. The writing on steering reminds me of one of my favorite jokes:

A caterpillar watched a centipede moving through the rain forest (?) and was amazed at how he used all of his hundred feet in such a coordinated way. He went over and asked: "how do you do it? How do you walk with all of those feet?"

The centipede stopped and thought. He wanted to answer the caterpillar's question.

Unfortunately, he could not. Nor could he ever walk again.

Eileen Prince Northeastern University

Date: Fri Aug 28, 1992 6:21 am PST
Subject: lcs ii & pct ed

is Living Control Systems II a sequel to Blayliss' (1969?) title, or to another book of the same title generated by this group?

is this subject (pct) formally taught anywhere? perhaps by someone in this group? or is this group, as i perceive, the only source for interactive instruction? reading and experimentation are part of the learning process (and i am actively and eagerly engaged in both), but one also requires a place where naive, perhaps even stupid questions may be taken seriously if they are offered in good faith. this is not meant as a criticism of most of the activity here... the responsiveness of this group is by far the best among the several which i monitor. the question remains, however; by which channels is this discipline currently or expected to be carried on?

-----------< Cognitive Dissonance is a 20th Century Art Form >-----------
Eric Harnden (Ronin)
<HARNDEN@AUVM.BITNET> or <HARNDEN@AMERICAN.EDU>
The American University Physics Dept.
4400 Mass. Ave. NW, Washington, DC, 20016-8058
(202) 885-2748
---------------------< Join the Cognitive Dissidents >-------------------

Date: Fri Aug 28, 1992 8:35 am PST
Subject: Influence, centipedes, instruction

[From Bill Powers (920828.0900)] Greg Williams (920828)--

>Once again, Skinner's "control" = PCT's "influence" -- consistently. If
>"most people think about control just the way Skinner did," then they
>believe that people can be influenced and that they themselves can
>influence people to have particular outputs. And that belief IS in fact
>true, isn't it, at least in a circumscribed way?

The contradiction here is contained in "influence people to have particular outputs." If you are only _influencing_ behavior -- which means that you can alter only one of many sources of influence and cannot determine whether there will be any effect -- then you can't expect to make a person produce particular outputs. The outputs will be whatever is required by the sum of ALL influences, not just yours, and it will also be whatever is required by the reference signal defining the desired outcome of the actions. Not the outcome you desire; the outcome the behaving system desires.

People are not generally content merely to apply an influence to someone else's behavior and accept the result. To do that would be to accept the fact that the other's behavior is most likely to proceed as before with no effect from the supposed "influence." What happens in reality is that if the first attempt at influence fails, as it is most likely to do, there will be continued attempts involving varieties of influences and increasing force behind the influences. The influencer reveals the fact that this is an attempt at control, not merely at influence. The object is to have a particular effect on the other's behavior that matches the influencer's goal for the effect.

The transition from mild and innocuous attempts at influence to concerted attempts to control is inevitable, for the simple reason that mere influence has almost no effect. The father, full of good nondirective and noncoercive intentions, says, "Johnny, don't you think you'd be safer if you didn't play in the street?" Johnny thinks it over for 0.5 sec and goes on bouncing his ball in the street. Papa gets up, goes to Johnny and looms over him and says "Really, is this a good place to play? Wouldn't you like to play with your ball on the nice soft lawn?" Johnny, not yet getting the point, says "It won't bounce on the grass." Papa ever so gently puts his arm around Johnny and pushes a little toward the curb, then a lot, and finally picks up Johnny, who is now struggling and screaming, and carries him to safety, putting him down on the nice soft lawn with perhaps a little extra velocity and handing him the ball. Papa says, "There, now, isn't this better?"

Papa knows where he wants Johnny to be, and puts him there after first trying some ineffective influencing. He's madly pretending, of course, that he isn't trying to control Johnny; that's why he only asks questions and speaks as if he's appealing to what Johnny really, deep down, wants. That's a delusion: he's talking about what HE wants Johnny to want. He can't go around following Johnny all the time; he's hoping that nice words will make Johnny do what Papa wants without Papa having to get up out of his nice lawn chair.

The delusion behind influencing people's behavior is the conviction that it works. By the time the "influencing" has finally been made to work, it has escalated into out-and-out control. Sweet reason finally comes down to "Get in the goddamned car; I'm not leaving a fertile 14- year-old girl alone in the house for two weeks." Although perhaps "fertile" would not be mentioned.

Psychology has always been concerned with the control of human behavior. It used to say so right out loud: the goal of psychology is the prediction and control of behavior. At the same time, however, the hope has been to accomplish this end through words, through little rewards, and if not punishments, at least withholdings and time-outs. In other words, to achieve control without all the nastiness and inconvenience of physically overpowering the victim. To do it by remote control, from the lawn chair. That's why the idea of "influence" is so appealing. You can't accuse someone of brutality just for saying "Wouldn't it be nicer to play on the lawn?" Once you get that sort of mild statement to work, however, you've already put the influencee on notice that if that doesn't work, there's more direct action that one is willing and able to take. What seems to be mere influence is understood by all concerned, whether said aloud or not, to be a warning that control is going to succeed, whatever it takes. So Johnny walks sullenly back to the lawn while Papa sits at ease sucking contentedly on his beer. Johnny then throws his ball onto the roof.

If you're going to control someone, then control -- don't pretend something else is happening. The pretense just leads to a psychotic society in which everyone talks and thinks as if one kind of relationship were going on, while the actuality is something entirely different. If we don't recognize when we're controlling, we will never figure out how our own attempts to control others are causing all the social ills we're trying to cure through control. We'll shy away from necessary acts of control, as in getting Johnny the hell out of the traffic, when control is actually called for. As long as we use euphemisms like "influence" when we really mean "control," we'll never discover how influence can actually be created without any tinge of coercion. And we'll go into conflict with ourselves when control is actually required.

Once we understand the difference between influence and control, we'll present whatever influences we think may help others or ourselves and be content if others decide that this time they choose not to be influenced. When we need to control someone to preserve a life or protect against being controlled, we will then not hesitate; we will do it quickly and skillfully using no more effort than necessary, and no less.

-- -

Eileen Prince (920828) --

RE: the centipede who couldn't walk. Gee, I hope I haven't caused a lot of accidents!

-- -

Eric Harnden (920828)

Living Control Systems (published works) and Living Control Systems II (unpublished ones) wasn't explicitly named for Bayliss's book, but we knew of it. Sort of a nod to Bayliss (sp?). He wrote a pretty good book.

Sources of instruction are a continuing problem. Until we get real CT courses in more colleges and universities, about all there is is reading and the net.

Don't, by the way, worry about asking stupid questions about simple things. And if you don't understand the answer, say so and we'll try again. Most of us on this net have the goal of helping people to understand control theory, not of proving how advanced we are. Often the simplest questions lead to the most interesting interchanges and encourage silent listeners to get involved in the conversation.

Best to all, Bill P.

Date: Fri Aug 28, 1992 8:56 am PST
Subject: Re: steering

[Ray Allis 920828.0845]

I memorized this from a "Sugar Creek Gang" book about 40? years ago:

A centipede was happy, quite,
until a frog, in fun,
said "Pray, which leg comes after which?"
This roused his mind to such a pitch
he lay distracted in the ditch
considering how to run!

It helped make me aware that there's a lot more to intelligent behavior than conscious, logical "reasoning". "AI" as presently constituted is a crock!

Ray Allis

Date: Fri Aug 28, 1992 9:54 am PST
Subject: Re: steering

Thank you. I have never seen the poem before. And I agree with the first sentence of your comment. I haven't the expertise to comment on the second.

Eileen Princu
Nortxeastern University

Date: Fri Aug 28, 1992 10:06 am PST
Subject: Paradox of Control

[From Rick Marken (920892.1000)]

I started writing this reply to Greg before I read Bill's wonderful reply this morning. I'll just post what I had written up to that point to prove that great minds think alike:

Greg Williams takes issue with my interpretation of Skinner's ideas about control:

Greg says (after an appropos Skinner quote):

>Suffice it to say, once again, that Skinner did NOT "use 'control' in the PCT
>sense" AT ALL. Q.E.D. (or R.I.P., or maybe not so peacefully, since the PCT
>bulldozer is headed Skinner's way!)
>I think you need to read what the opposition (loyal or not) says before
>putting words into their mouths. Your turn to quote scripture.

I'm willing to believe that Skinner and all the others who advocate behavior control as a solution to behavior problems will be able to verbally message the meaning of "control" as necessary so that it doen't mean "has the purpose of producing a particular behavior". When I talk about the paradox of control I'm talking about what is actually going on when one person tries to get another person to perform a PARTICUALR behavior. This person is claiming is trying to control behavior (in the PCT sense) whether he will admit it in print or not. So, no matter what Skinner says (and he says plenty) he advocated controlling behavior while claiming that he could do this because all behavior is organized in a way that would obviate his own ability to exercise control. Whether he would admit it or not, Skinner was caught squarely in the paradox of control. Poor B.F. was a control freak who simply refused to try to undertsand what he was so excited about -- CONTROL.

eric harnden asks:

>is this subject (pct) formally taught anywhere?

Good question. I think Dick Robertson teaches it in Chicago. I gave a seminar in it several years ago when I was a professor. Tom Bourbon taught it in Texas. Chuck Tucker, Gary Cziko, Clark McPhail and Wayne Hershberger and Kent McClelland also come to mind as people who might have given courses in it (they are real professors still). Ed Ford teaches it for clinical audiences, Dag Forssell teaches it in industry. But basically CSG-L is the main forum for learning PCT right now. PCT is certainly not part of the standard psychology curriculum yet -- indeed, if people ever got it the PCT would BE the standard psychology curriculum, which is one reason PCT will probably not be showing up in many classrooms soon.

> one also requires a place where naive, perhaps even stupid questions
>may be taken seriously if they are offered in good faith.

I hope that this net is just such a place. I certainly welcome questions and I certainly don't consider ANY question stupid. There is no contest on CSG-L to prove who's the smartest, quickest, or whatever. We are ALL here to learn and teach all we can about an extraordianary insight into the nature of living systems known as PCT. I hope that my own rants have not scared you. I am the veteran of many years of trying to present PCT to an audience that has absolutely no interest in learning it (scientific psychologists, mainly, and the reviewers who stand at the gates of their journals). But I'm really quite tame -- and I am quite sure everyone is a hell of a lot smarter than I am, so I have nothing to prove. I do understand and love one thing (besides the wife and kids) -- PCT. So PLEASE feel free and comfortable to ask and challenge; we need it.

Best regards Rick

Date: Fri Aug 28, 1992 10:08 am PST
Subject: candidate models of reorganization

[From: Bruce Nevin (Fri 920828 13:21:43)]

It doesn't take long to fall far, far behind the stream here. I will be gone all of next month, so I suppose I will just have to tune in as a newcomer when I get back :-) But first, one note out of many scribbled while on the train.

(Martin Taylor 920820 19:15) --

> I think we have three candidate reorganization systems that are plausible,

Add to this one more, Bill's fallback position and the one I like as primary, though it is harder to model. In this view, reorganization is carried out by the constituents (cells) of the control system experiencing conflict and chronic error. Call the latter a control system of order n; then its cells or constituents are a population of control systems of order n-1 (e.g. nerve cells). The hypothesis is that when a control system of order n experiences error, it produces something intrasomatically that, if it persists, is the occasion of error in its cells as control systems of order n-1. They then begin to change shape, contacts with one another, reactiveness, etc. They do so until their local error is reduced. But reducing that order-n-1 error corresponds to reducing the order-n error, in precisely the sense that reorganization continues until it reduces the error that occasioned it.

This is difficult to model. Would you have soldered joints respond to overheating by rearranging themselves? The problem, of course, is that soldered joints are not control systems. (Similarly for other constituents of our means for modelling). But conceptually, intuitively, this is very satisfying to me. Seems to fit with Edelman's ideas too.

Bruce bn@bbn.com

Date: Fri Aug 28, 1992 11:57 am PST
SUBJECT: PCT in the Classroom

[FROM Dennis Delprato (920828)]

I supervise laboratory experiences of approx. 380 1-2-year students and 60 3-4-yr. psychology majors each year. For this I have worked up numerous labs that have several administrative requirements. Basically, they are prepared such lab instructors can conduct the labs with minimal direction from me. Each lab has a manual of instructions with equipment, materials, procedural, and such requirements spelled out for the instructors. Each also has a handout for the students for them to use in the lab meeting, for reference, and for lab reports, when required. I lay out the above by way of seeking assistance from anyone who might be interested in getting a PCT lab back on line for use here and anywhere else (I am not interested in copywriting the thing). I am interested in getting the lab "up."

As the late great Ross says, "Now, here's the story":

I do no programming, but was able to get the assistance of a capable student who just about got a particular PCT lab on line. I based the lab on Rick Marken's "Selection of Consequences" paper (Psychol. Reports, 1985, 56, pp. 379-383 also see Marken & Powers (1989, Behav. Neuroscience, 103, 1348-1355) The lab is entitled "Goal-Seeking with Random Consequences of Responses."

I require my labs to go as smoothly as possible from beginning to end. In the case of this lab, this means the student must have guidance provided from the time the Macintosh is turned on until they turn it off. No problem. This goal was met. I like flexible labs, by which I mean variables can be changed. No problem. We have default values for several variables, but they can be changed. I like a clear demonstration for students. No problem. The E. coli problem that Rick devised is a beaut. I made it into a game. I require data to be collected. No problem. The computer stores data and the student can print it out at the end. Labs that "involve" students because they are "interesting" are a plus. No problem here.

Given the lack of problems to this point, then just what is "my" problem. The problem is that the "E. coli" program causes the Macs to crash in an unpredictable (more or less) fashion. Not only does the program crash, but other files on the hard drive get contaminated. Now I greatly admire what my programmer has done to this point, but working with amateurs can be frustrating. The first things we looked for were viruses. No matter how hard we looked, we could find no evidence of viruses. The last hypothesis had something to do with the fact that the program used Rascal and perhaps a Mac operating system that has some slight difference from the ones in the Macs here. The student was at Mich. State when he did the programming. It could be the Rascal. It could be the operating systems, ...

Just to fill in on the task, the "EColi Game" is introduced to the student by directions on the screen that read as follows:

"Your goal is to move the cursor into the target [a box with a size that can be varied] as fast as possible. There will be a time limit of five minutes."

"The only means that you have to control the cursor is to click the mouse button in the Executor Window. Moving the mouse has no effect on the cursor."

"When the game ends, you time and some other relevant data are shown."

"If you need to stop the game before five minutes have elapsed, click the mouse inside the target area." [I recall noting that some other way is needed for this.]

"Hit any key to continue."

IN ANOTHER VERSION, instead of measuring the time to reach the target, a run (trial) is set at a partucular duration and the student earns points derived from the amount of time the cursor stays inside the goal (target, box).

Per Rick Marken's E. coli preparation, responses (mouse clicks) move the direction of the cursor according to a random pattern.

Each student can collect a nice amount of data rather quickly (we have 6 or so Macs). Performance over runs (trials) can be plotted and even GROUP data (I know, shame) can be analyzed via repeated-measures AOV. The latter can be used to show how uninformative is the analysis of group data in that wide individual differences are revealed in the data.

Students can even print out a graphic of their cursor's tracks.

As I have implied, I have been able to have some computers go through the entire data collection process. It is just that the crashing is so unpredictable to render routine running of the lab unwise.

I suppose I have supplied enough information to anyone who might be interested. The recent comments on PCT in the classroom led me to think of this undertaking that I have let sit for about a year.

Dennis Delprato
Dept. of Psychology
Eastern Mich. Univ.
Ypsilanti, MI 48197 U.S.A.
Psy_Delprato.EMUNIX.EMICH.EDU

Date: Fri Aug 28, 1992 11:58 am PST
Subject: A General Theory of Manipulation

From Greg Williams (920828 - 2)

>Bill Powers replying on "influence" (920828)

OK, Bill, I could counter your lurid family-relations examples showing the ineffectuality of influence with lurid con-game and advertising examples showing the effectuality of influence, but what's the use? Instead, why not join with me in constructing a general treatment of manipulation from the point of view of PCT -- a general treatment which will show the scope and limits of manipulation? Here's a start. The axioms are supposed to arise from orthodox PCT.

PROLEGOMENON TO A GENERAL THEORY OF MANIPULATION

Definition 1. Control is acting so as to maintain a match between reference signals and perceptual signals in the face of disturbances which tend to alter the perceptual signals.

Axiom 1. Comparing reference signals and perceptual signals occurs ONLY within organisms, including humans.

Axiom 2. Given current and forseeable technology, other persons CANNOT directly ("causally") alter a given person's current structure of reference signals.

Axiom 3. Other persons CAN, in general, directly ("causally") generate disturbances which alter a given person's current perceptual signals.

Lemma 1. A person can ONLY control his/her own actions, not others. In particular, other persons cannot control a given person's actions. (By Def. 1 and Axioms 1-3.)

Definition 2. Influence is one person's controlling his/her actions so as to alter another person's perceptual signals.

Lemma 2. Influence is possible. (By Def. 1 and Axiom 3.)

Definition 2. Manipulation is a given person's (the influencer's) altering of another person's (the influencee's) current perceptual signals so as to affect the influencee's current control processes in such a way that the influencee acts so as to enable matching of certain reference signals and perceptual signals of the influencer (that is, to enable certain control processes of the influencer).

Lemma 3. Manipulation is possible. (By Defs. 1 and 2 and Axioms 1, 2, and 3.)

Axiom 4. Information about the influencee's current control structure (that is, structure of reference signals) cannot be obtained directly, given current and forseeable technology, but can be obtained indirectly, by the Test for the Controlled Variable.

Axiom 5. Successful manipulation as set forth in Def. 2 would require information about the influencee's current control structure so as to avoid conflicts within the control structure and to assure consonance between the current control processes and acts of the influencee which enable control processes of the influencer.

Theorem. By Lemma 3 and Axioms 4 and 5, the chances of successful manipulation are enhanced IF the influencer performs the Test.

This would be extended in a full treatment to mapping out how and why manipulation has limits, the possible manipulation of others' FUTURE control structures by altering their PRESENT perceptual signals, and so forth. Are you game?

Greg

Date: Fri Aug 28, 1992 11:59 am PST
Subject: Re: candidate models of reorganization

I am fairly new on this net so I hope this is not a "could have been answered if you looked here" question ... but, do any of the models predict the top layer to change and if not how are the top levels set in the first place? If the top levels change (and intuitivly they should), do they change themselves, help change themselves, or are they changed by the lower levers to create a sort of circle? Also, how are connections between levels set in the first place, espicially the top level as it has the potential for being the most complex wrt sensory inputs?

-John van Loon
-no fixed address, anymore

Well I have to delete my desk top now so I won`t get the answers I want untill I sign on again. Thanks for controlling my thoughts this summer, course I can`t drive anymore ... Oh well I think that I can still walk.

Date: Fri Aug 28, 1992 12:25 pm PST
Subject: One more GTM axiom

From Greg Williams (920828 - 3)

I forgot an axiom in my last post, as Pat just reminded me:

Axiom 6. PCT has nothing to say about the utility of attempted manipulation, in general, for either influencers or influencees, except that it might waste time and/or resources for either or both parties.

Greg

Date: Fri Aug 28, 1992 1:43 pm PST
Subject: PCT in the Classroom

[From Rick Marken (920828.1330)]

Dennis Delprato (920828)

>I supervise laboratory experiences of approx. 380 1-2-year
>students and 60 3-4-yr. psychology majors each year.

Sorry, forgot include your name in the list of PCT teachers.

What kind of Macs do you have? Do you have Hypercard 1.0 or 2.0? Which one?

I'll try to write a program that will support your lab. I think I could do it in Hypercard -- but the animation might not run fast enough unless you're running an SE30 or better. I could also write it in Basic or Pascal. In either case I could send you both the source and object files. This might be the best way to go if you have Turbo pascal or microsoft basic.

Let me know. I'd really like to do it.

Best regards Rick

Date: Fri Aug 28, 1992 8:04 pm PST
Subject: Marken Computer Demos

[from Gary Cziko 920828.2237]

Rick Marken (920828.1330) to Dennis Delprato (920828)

>I'll try to write a program that will support your lab.
>I think I could do it in Hypercard -- but the animation might
>not run fast enough unless you're running an SE30 or better.
>I could also write it in Basic or Pascal. In either case I
>could send you both the source and object files. This might
>be the best way to go if you have Turbo pascal or microsoft basic.

Rick, some time ago you mentioned the possibility of "packaging" your programs into one with a menu interface. I would find this very useful for my students. We do use programs like Find Mind and Mind Reader, but having it all together would be a lot easier. Would this fit into Dennis's needs as well?--Gary

Date: Fri Aug 28, 1992 8:04 pm PST
Subject: Prolegomenon; Misc

[From Bill Powers (920828.1800)] Greg Williams (920828-2) --

RE: Prolegomenon. I didn't know that Machiavelli was Greek.

>... why not join with me in constructing a general treatment of
>manipulation from the point of view of PCT -- a general treatment which
>will show the scope and limits of manipulation?

While we're at it, let's be sure to examine hidden assumptions.

Definition 1 and Axiom 1 state that control consists of comparing perceptual signals with reference signals and acting to maintain them in a match. Axiom 2 states that people can't alter each others' reference signals. The definition and the two axioms thus indicate that there is no way for people to alter each others' perceptual signals if those signals are under control. We exclude, I assume, the application of overwhelming physical force as being outside the meaning of "manipulation."

This conflicts with:

>Axiom 3. Other persons CAN, in general, directly ("causally") generate
>disturbances which alter a given person's current perceptual signals.

The only way to resolve this conflict is to insert the word "uncontrolled" between "current" and "perceptual signals."

Disturbances in the sense of Axiom 3 (distinguished from any alterations they may create) are independent variables under the control of the disturber. They must be elastically linked to the perceptions they affect if those perceptions are under control; otherwise we have a case of simple conflict, not manipulation. In speaking of either disturbances or influences, we must distinguish between the independent variable (the potential cause of a change) and its actual effects at the other end of the link.

The same problem arises with:

>Definition 2. Influence is one person's controlling his/her actions so
>as to alter another person's perceptual signals.

There will be an alteration, by Definition 1 and Axioms 1 and 2, only if the perceptual signals in question are not under control. Thus we need to change Definition 2 to end with "uncontrolled perceptual signals."

We now have a non-sequitur in Definition 3 (mislabelled as a repeated Definition 2):

Definition [3]. Manipulation is a given person's (the influencer's) altering of another person's (the influencee's) current perceptual signals so as to affect the influencee's current control processes in such a way that the influencee acts so as to enable matching of certain reference signals and perceptual signals of the influencer (that is, to enable certain control processes of the influencer).

Altering the influencee's UNCONTROLLED perceptual signals will not affect the influencee's control processes. By Definition 1 and the first two Axioms, controlled perceptual signals cannot be altered (without escalating the situation beyond the definition of "manipulation"). So if manipulation occurs, it does not occur through alteration of the manipulee's controlled perceptions. It can occur only through altering uncontrolled perceptions.

This brings us to The Test.

>Axiom 4. Information about the influencee's current control structure
>(that is, structure of reference signals) cannot be obtained directly,
>given current and forseeable technology, but can be obtained
>indirectly, by the Test for the Controlled Variable.

The Test reveals not only the reference signal, but the perceptual signal that is under control. At the same time, it identifies something that is, at the level of concern, NOT under control: the action taken by the person. The disturbances used during the Test succeed in altering the action, but not the perception. It is the failure to alter the perception that reveals the controlled perception. So the Test, aside from identifying reference levels and controlled perceptions, also shows what variables the person will freely alter as necessary. This identifies the variables that are, in the context, of no importance to the person.

We thus arrive at a conclusion: it is possible for one person to manipulate aspects of another person's behavior as long as those aspects are of no importance to that other person. The Test does, therefore, enhance the chances of successful manipulation, as your Theorem proposes. However, it does not enhance the chances of manipulating anything of importance to the manipulee. The Test identifies those behavioral variables that can be manipulated (actions) and those that cannot (controlled perceptions).

I think that the Definitions and Axioms need to be modified.
 -- --
Going back to the opening of your post:

>OK, Bill, I could counter your lurid family-relations examples showing
>the ineffectuality of influence with lurid con-game and advertising
>examples showing the effectuality of influence, but what's the use?

The use would be to take these counterexamples apart and show that in every case they require the active cooperation of the manipulee. Consider a con man, and the Pigeon Drop.

The con man says that he has just found this folder full of bearer bonds, but he can't cash them in because he's adversely known at the bank. Will you do it for him? Of course you have to put up some security before he hands over the bonds. At the bank, you find that the folder is full of newspaper. Tsk, tsk.

He manipulated your perceptions, right? No. YOU manipulated your perceptions; you're the only one who can.

Now the con man goes to his next victim. The next victim listens carefully, and calls a cop. What made the difference?

The difference was in the manipulees, not in the manipulator. The first guy could have been just as skeptical -- but the first guy wanted to get something for nothing, and simply refused to consider any information that would disappoint him. The first guy played the con man's game with him in order to get something he wanted. The con man makes a living by playing this game with people who want to play it. The con man can cheat poor little old ladies out of their life savings not because he can manipulate them, but because he looks for greedy little old ladies who are only too ready to play the game. That is one reason that they are POOR little old ladies.

Now this may sound like blaming the victim, and in a way it is just that. I am not very interested in developing a manual for manipulators. I'm interested in teaching people how to avoid being manipulated, which would obsolete any such manual. I can't teach them to do this by showing how to make con men change their ways. If your reference signal whispers "Get Rich Quick" in your mental ear, nobody else in the world can change that, and you will suffer the consequences that you make for yourself all by yourself, even without the help of a con man. What you have to learn is that manipulation can't work on you unless you actively help it along.

A very long time ago, the Buddha put his finger on the problem: it's the wheel of desire. Whatever you want determines what you will do to get it. The more intensely and inflexibly you desire something, the less control you have over your own actions. The wheel of desire looks something like the CSG logo.

John van Loon (920828)

The top level obviously can't change its own reference signals; they are inputs to it, not outputs from it. If they change we have to look elsewhere for the mechanism. Knowing as little as we do, all we can do is entertain possibilities.

One possible mechanism is memory. What you have experienced at this level becomes the reference-state of perceptions at this level. This has some explanatory power: it suggests that the world we live in becomes the status quo, and then the desired status quo.

Another is reorganization. The highest levels of reference signal are picked at random, the one you finally settle on being the one that leads to the least intrinsic error in your reorganizing system. This, too, has some explanatory power, particularly with respect to how "upbringing" leads you to adopt system concepts that lead to the least pain, rejection, and conflict.

Another -- which works only for some system concepts -- is to say that there are no reference signals at the highest level, which is equivalent to saying that they are set to zero. This explains the automatic rejection of system concepts with which one has had little experience.

Another is to say that God sets your highest reference signals, so if they don't work out it's not your fault.

In the meantime, there are several problems we could work on at the lower levels while we wait for the explanation of the highest level to become clear.

By the way, have you ever looked carefully at how you get from sitting down to standing up balanced directly over your feet?

-
Rick Marken (920828) --

While I was reading Dennis Delprato's post I was wishing you would volunteer your Mac expertise. Bully for you.

Best,   Bill P.

Date: Sat Aug 29, 1992 4:19 am PST
Subject: Unhidden assumptions in GTM

From Greg Williams (920829)

>Bill Powers (920828.1800)

>RE: Prolegomenon. I didn't know that Machiavelli was Greek.

Interesting you should say that. It turns out that people of Greek descent (also Chinese) are considered hard to sell by (small- or big-) con artists. Too "honest"??

GW>>... why not join with me in constructing a general treatment of
GW>>manipulation from the point of view of PCT -- a general treatment which
GW>>will show the scope and limits of manipulation?

BP>While we're at it, let's be sure to examine hidden assumptions.

>Definition 1 and Axiom 1 state that control consists of comparing
>perceptual signals with reference signals and acting to maintain them
>in a match. Axiom 2 states that people can't alter each others'
>reference signals. The definition and the two axioms thus indicate
>that there is no way for people to alter each others' perceptual
>signals if those signals are under control.

True, but I can provide the opportunity for you to control perceptual signals which you would certainly not have been controlling had I not been around ("Come over and look at this wallet I found here!"). I cannot be sure that you will come over, of course, because that depends on your current control structure (reference signals) -- but without my "disturbance" (ALTERING your perceptual signals, as I said in the "Prolegomenon," but perhaps more to the point, leading to an alteration in WHICH perceptual signals you are controlling), you would certainly NEVER come over and look, since there would be nothing to look at. I think this restatement obviates your critique.

>The use would be to take these counterexamples apart and show that in
>every case they require the active cooperation of the manipulee.

This was exactly my point. PCT says that for manipulation to have a chance of success (setting aside coincidence/"lucky" guessing by the influencer), the influencer must try to alter the influencee's perceptual signals so that the influencee's control processes will result in certain perceptions of the influencer. You and I are in complete agreement on this.

>He [a con man] manipulated your perceptions, right? No. YOU manipulated your
>perceptions; you're the only one who can.

In general, the con men alter the marks' "don't-care" perceptions in such a way as to SHIFT the marks' controlling processes so as to result in certain perceptions of the con men. This is perhaps a clearer expression of what I was trying to express in the "Prolegomenon." Again, you and I are in complete agreement (unless you hold that one individual's disturbances of another's "don't-care" (uncontrolled) perceptions can never result in the latter's shifting to control perceptions which he/she otherwise would not have been controlling). Your (earlier, above) bringing up the importance of "don't-care" perceptions in the manipulation process has provided me with the opportunity (which I otherwise would not have had) to make this PCT-based theory more detailed (which I want to do). Thanks for the manipulation! You get the "yes, Bill, you are right" perception.

>Now this may sound like blaming the victim, and in a way it is just
>that. I am not very interested in developing a manual for
>manipulators. I'm interested in teaching people how to avoid being
>manipulated, which would obsolete any such manual.

Who said that manipulating or being manipulated is always (or usually) bad? I think it is part and parcel of everyday life, and has both good and bad effects for influencers and influencees in particular cases. I don't think it would be a very pleasant world if everyone steadfastly refused to manipulate or be manipulated -- have you ever lived around a "paranoid"? I do think that PCT can provide the basis for a general theory of the nature and limits of manipulation, giving insights into its everyday operations, and possibly usefully exposing con games.

>What you have to learn is that manipulation can't work on you unless you
>actively help it along.

PCT can help to flesh out assertions like this. Again, are you game?

Greg

Date: Sat Aug 29, 1992 7:48 am PST

From: Dag Forssell / MCI ID: 474-2580
TO: list (Ems)

set csg-l repro

Date: Sat Aug 29, 1992 7:56 am PST
From: Revised List Processor
EMS: INTERNET / MCI ID: 376-5414
MBX: LISTSERV@vmd.cso.uiuc.edu

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Output of your job "0004742580"

> set csg-l repro
Your distribution options for list CSG-L have been successfully updated.

Date: Sat Aug 29, 1992 11:16 am PST
Subject: Re: candidate models of reorganization

[Martin Taylor 920829 13:00] (Bruce Nevin 920828 13:21:43)

>(Martin Taylor 920820 19:15) --
>
>> I think we have three candidate reorganization systems that are plausible,
>
>Add to this one more, Bill's fallback position and the one I like as
>primary, though it is harder to model. In this view, reorganization is
>carried out by the constituents (cells) of the control system
>experiencing conflict and chronic error.

So far, it sounds identical to my own position, at least if by "control system" you mean "ECS". If you mean "hierarchy" I don't see how a hierarchy experiences anything, conflict and error included. An ECS can experience error, which could be the result of conflict, but it can't experience conflict directly. You go on to expand on how reorganization might be achieved in a specific kind of hardware(mushware) implementation. I have no objection to that, but the way it is done would have to be different in other (say intracellular) systems. I am afraid I have never properly understood your "control system of order n" concept, so maybe I am quite misreading your intention here. My notion works entirely within a (any) single control hierarchy, Bill's requires a separate hierarchy monitoring variables that do not appear in the percepts of the hierarchy being reorganized, and Jeff's requires a set of symbiotic control systems that may be hierarchies, but that influence the one being reorganized only by inducing error into it so that it reorganized according to my (ad your?) scheme. Does your "order n" and "order n-1" fit into one of these three categories, or is it something else again?

Martin

PS. Personally, I rather like Jeff's version, although Bill thinks that it can't work. In time, I hope, we shall see.

Date: Sat Aug 29, 1992 11:16 am PST
Subject: Re: A General Theory of Manipulation

[Martin Taylor 920829 13:15] (Greg Williams 920828 - 2)

I think a lot of the problem in the discussion of "controlling" other people comes from a dual use of the word "control," which was much clarified by Bill's posting of yesterday (reference not at hand). Using Greg's set of axioms indirectly, or Bill's discussion, we have a statement that the only legitimate use of "control" is the control of perception, specifically the provision by an ECS of output signals that by some unknown means bring the percept of that ECS nearer to the reference level of that ECS. Any other use of the word "control" is analogic and imprecise. Any percept in an ECS corresponds to a CEV (complex environmental variable). It is an estimate, more or less precise, of the current state of that CEV. The word "control" is often used to indicate effects on the CEV, but that is only legitimated by the assumption that the percept is a precise and immediate representation of the exact state of the CEV. One should never say that one "controls" the location of the car within its lane on the road. One should say that one controls the percept of the location of the car.... The word that has been used to describe external effects has been "influence," but this seems inappropriate, because it has the S-R connotation that a force has been applied, but the result not noted. We need a different word for what happens to the CEV corresponding to a controlled percept. For lack of a better word, I propose here to use X-control (externalized control). One can X-control passive CEV's in the real outer world. One cannot control (P-control?) them. Bill pointed out that one does incorporate the actions of other people in control loops. We want the child to be in the garden, not in the middle of the busy street. If one action does not bring this desired percept to its reference, another might, and we continue acting until the percept is within its reference. We have controlled the percept of the child's location. We have X-controlled the child's location. We may have done this by restricting the child's range over which it is capable of controlling its own perception of its own location, or by influencing it to change its reference level for its percept of its own location. Whatever happened, we X-controlled the child in at least one aspect. From our point of view, it may be more difficult to X-control a child than to X-control a car, but the principle is the same.

I hope this clarifies something.

Martin

Date: Sat Aug 29, 1992 11:16 am PST
Subject: Re: Prolegomenon; Misc

[Martin Taylor 920829 14:45] (Bill Powers 920828.1800)

>Definition 1 and Axiom 1 state that control consists of comparing
>perceptual signals with reference signals and acting to maintain them
>in a match. Axiom 2 states that people can't alter each others'
>reference signals. The definition and the two axioms thus indicate
>that there is no way for people to alter each others' perceptual

>signals if those signals are under control. We exclude, I assume, the
>application of overwhelming physical force as being outside the
>meaning of "manipulation."
>
>This conflicts with:
>
>>Axiom 3. Other persons CAN, in general, directly ("causally")
>generate
>>disturbances which alter a given person's current perceptual signals.
>
>The only way to resolve this conflict is to insert the word
>"uncontrolled" between "current" and "perceptual signals."
>

Wait a minute here...This is one of the things I was talking about when I said that the dynamics are important. You are dealing with steady states, whereas the effect the attempting controller would have is transient. Even in the steady state, when there is conflict, two controllers attempting to move a single CEV to different levels, each will show some persistent error, so even in the steady state one can influence the other's controlled percepts. How much the "controller" influences the "controlee"s percept will depend on who has the higher insistence (gain), and that could be interpreted as who applies the more overwhelming physical force. We do not exclude that, I think, because it applies to a greater or lesser degree in every instance of conflict.

Besides, does planting a fence constitute "overwhelming physical force?" It certainly affects the child's ability to perceive itself as being in the middle of a busy street. The parent's actions do affect the child's controlled percepts, even in the steady state and without the parent experiencing conflict.

Martin

Date: Sat Aug 29, 1992 1:41 pm PST
Subject: sonja's visuals system

[Avery Andrews 920830.0735]
Bill Powers (920827.0800)

>can you explain how Sonja's visual system identifies the various
>retinal images, and how the relative retinal distances between image
>points are converted into signals?

Not as well as Chapman can: although it would an excellent exercise for me to try, I think I've got too much else on my plate at the moment to take up the challenge.

Avery.Andrews@anu.edu.au

Date: Sat Aug 29, 1992 9:19 pm PST
Subject: Manipulation

[From Bill Powers (920829.2130)]

RE: various posts on the subject of manipulation.

The context of Greg Williams' discussion of manipulation requires a particular interpretation of the term. What is meant is one person controlling another person's actions surreptitiously, leading to that other person doing something one could not get the person to do simply by asking. The implication is that the other is made to behave in some way independent of his or her interests or wishes, without arousing a direct conflict or direct resistance by the other person. The means of manipulation are not the application of direct force, or a request for the other to perform the act, but deceit, indirection, misrepresentation of facts, false promises, ambiguous communications, and (using the Test) disturbance of variables the victim is known to be controlling in order to elicit an opposing response with consequences unknown to or unpredicted by the victim. While manipulation in this sense might be used for the benefit of the manipulee, that is a judgement of the manipulator, not the manipulee; it is the goals of the manipulator that are satisfied, while those of the manipulee are ignored or bypassed.

Greg's thesis is that there are ways of controlling the behavior of other people against their will or without their awareness other than by the use of overwhelming physical force.

My thesis is that such apparent control (a) is an illusion, or (b) is of no importance to the "victim", or (c) is carried out only with the aid of the "victim," or (d) rests on a threat of force in the background. I contend that a properly functioning control system cannot be controlled from outside, so that any apparent success by a manipulator results from some defect of organization in the control system.

Martin Taylor (920829.1315) --

>One should never say that one "controls" the location of the car within
>its lane on the road. One should say that one controls the percept of the
>location of the car....

The stages of Satori. Before I understood, I said that I controlled the car within its lane. While I was understanding, I said that I controlled the percept of the car within its lane. After I understood, I said that I controlled the car within its lane (for I know only the percept in any case).

>We need a different word for what happens to the CEV corresponding to a
>controlled percept. For lack of a better word, I propose here to use
>X-control (externalized control). One can X-control passive CEV's in
>the real outer world. One cannot control (P-control?) them.

I resist this terminology. If the (hypothetical) external CEV (complex environmental variable) does not behave as the percept does, then the CEV has been misdefined; it must entail aspects that are not matched in the perceptual function of the controlling system. While I don't object to speaking of a model of the physical reality as if it really exists (practical considerations require this), I want to minimize reification of that hypothetical reality. To speak of a CEV apart from the perceptual system that defines it for a controlling system is to give the CEV an independent existence of its own, to objectify it, and to imply a third-party observer who can know the true form of the CEV which may be partly misrepresented in the controlling system. There may be many variables in the external world. But they are not organized into CEVs.

It is the perceptual function that defines the true relationship between the percept and that which is externally controlled. The percept can't be incorrect. Only our understanding of its external counterpart can be in error. That is, our assertion that we are perceiving the world in the same way that the control system is perceiving it is fallible. There are no CEVs independent of the controlling system (or the observer). A correctly-defined CEV always corresponds exactly to the percept.

There is only one form of control: control of the percept.

RE: manipulation of uncontrolled perceptions

>Wait a minute here...This is one of the things I was talking about when I
>said that the dynamics are important. You are dealing with steady states,
>whereas the effect the attempting controller would have is transient.

A disturbance that has a transient effect only will not materially affect the behavior (output) of the control system. If the control system could react within the interval we classify as "transient," the disturbance would be successfully resisted and there would be no transient effect on the perception. Remember that the main means of producing predictable outputs from a control system is to apply a disturbance THAT IS CANCELED BY THE ENSUING ACTION. If the disturbance is brief enough to qualify as a true transient that the control system cannot counteract before it is over, then there will be no action that is equal and opposite to the disturbance.

Control systems evolve and reorganize until they can react adequately on the time-scale typical of normal disturbances. Don't think of disturbances as transient "events." They are simply independent variables that influence the controlled variable. They may be brief, or like gravity they may be present permanently.

>Even in the steady state, when there is conflict, two controllers attempting
>to move a single CEV to different levels, each will show some persistent
error,
>so even in the steady state one can influence the other's controlled
percepts.
>How much the "controller" influences the "controlee"s percept will depend on
>who has the higher insistence (gain), and that could be interpreted as who
>applies the more overwhelming physical force.

I repeat that an influence (or disturbance) is neither control nor determination. It is simply an influence, which may or may not have any effect, depending on other influences that may be present at the same time (such as the output of a control system).

With two systems balanced in conflict as you describe above, each system is producing some amount of output that is canceling an equal amount of the output of the other system. Varying either reference signal will result in moving the balance point, but only until the effort of the active system reaches its limit. Beyond that point, neither system will have any control over the common variable. In control systems with integrating outputs (which seem to be common in real systems), both outputs will quickly reach saturation, and there will be no effect from varying either reference signal.

As I pointed out at the start of this post, "manipulation" in the sense that Greg Williams means does not include direct conflict, which can't be brought about surreptitiously.

>Besides, does planting a fence constitute "overwhelming physical force?"

If no matter how hard the control system tries it can't get past the fence, yes. The overwhelming physical force is automatically applied, to whatever extent is necessary, as soon as the control system sets the reference level for its position to the other side of the fence and attempts to correct the error. The fence totally overrides the organism's efforts at control. Control systems can be prevented from functioningd by overwhelming physical force, whatever the source of that force. Of course we would not say it is "controlled" unless the source of the force is intentional. It is not really the fence that controls the child's behavior, but the person who put the fence there.

Because the child can't control its position relative to a place outside the fence, it will reorganize so that it no longer desires to be outside the fence. It will then continue to live, to greater or lesser extent, as if it is inside a fence even when the fence is removed. One of the great difficulties in growing up is that of understanding which fences have been or should be removed, so there is no longer any need to restrict one's desires to the region within their confines.

Adolescents often rail against the constraints of "the system" (oh rueful memories), not yet realizing that the constraints are in themselves and that "the system" is largely a memory of the restrictions, punishments, and fears of childhood -- to which they adapted themselves all too expertly.

Best to all, Bill P. (Turned 66 today: happy birthday to me).

Date: Sun Aug 30, 1992 3:47 am PST
Subject: Happy Birthday, Bill (belatedly)

From Greg Williams (920830) >Bill Powers (920829.2130)

Happy Birthday + 1 day!!!!

>The context of Greg Williams' discussion of manipulation requires a
>particular interpretation of the term.

It certainly does. I defined it in a very particular, precise way.

>What is meant is one person controlling another person's actions
>surreptitiously, leading to that other person doing something one could not
>get the person to do simply by asking.

I thought we'd already reached an agreement that one person cannot control (in the PCT sense) another. Are you using "control" more colloquially here? If so, your statement is more or less OK, but it might be possible (at least in some cases) to achieve the "controller's" desires simply by asking. Still, I am concerned about your use of "surreptitiously"; my definition of manipulation did not specify that the influencer's desired perceptions be hidden from the influencee. An example is the diet center counselor who manipulates the dieter so that both influencer and influencee see fewer pounds on the influencee.

>The implication is that the other is made to behave in some way independent
>of his or her interests or wishes, without arousing a direct conflict or
>direct resistance by the other person.

Perhaps "in some way independent" at least some of the time, but what I was trying to get at was manipulation as making the influencee control (to SATISFY his/her interests or wishes) in such a way as to ALSO satisfy certain interests/wishes of the influencer. Yes, there should be no direct conflict within the control structure of the influencee, and no direct resistance by the influencee. This is why your previous comments about the influencer altering "don't-care" perceptions of the influencee is important.

>The means of manipulation are not the application of direct force, or a
>request for the other to perform the act, but deceit, indirection,
>misrepresentation of facts, false promises, ambiguous communications, and
>(using the Test) disturbance of variables the victim is known to be
>controlling in order to elicit an opposing response with consequences unknown
>to or unpredicted by the victim.

Some of the consequences could be unknown to or unpredicted by the influencee; other (or sometimes all) consequences (including the influencer's desired perceptions which are the "point" of the manipulation for the influencer) could be known to or predicted by the influencee. The latter possibility is in no way excluded by my definition of manipulation. This means that, according to my definition (which seems much broader than your own definition -- remember, I'm trying to make a folk notion more precise), misdirection need not always be "good" for the influencer and "bad" for the influencee. Also, your "... and (using the Test)..." should read "... and/or (using the Test)..." Really, that final clause was as far as my definition of misdirection went; you have been emphasizing only the, shall I say, NEGATIVE possibilities. Another example: Pat puts what she thinks is healthy food on the table; our kids (usually) gobble it up (they say it is really tasty!); influencer Pat's desired perception is seeing the kids eat "healthy" (her definition) food, and she manages to help the kids fulfill their desired perception of eating (their definition) "tasty" food. Now, if you ask one of our kids whether Pat is providing them with "tasty" food so that she sees them eating "healthy" food, they'll say, "Sure." There's no deceit, indirection, misrepresentation of facts, false promises, or ambiguous communications here, but there IS manipulation as I defined it, because Pat altered some of the kids' perceptual variables in such a way as to result in them being able to control for certain perceptions which also result in Pat being able to control for certain perceptions. Do you begin to grasp the pervasiveness of manipulation in everyday life?

>While manipulation in this sense might be used for the benefit of the
>manipulee, that is a judgement of the manipulator, not the manipulee; it is

>the goals of the manipulator that are satisfied, while those of the manipulee
>are ignored or bypassed.

No, certain goals of both parties must be satisfied for manipulation to work. Before a manipulation is successful, the influencer does indeed predict what the influencee will want, and could be wrong. But afterward, the influencee can judge for him/herself. In a con game, the influencee (sometimes) decides not to play again. In what might be termed "benign" or "positive" manipulation, the influencee would presumably decide to play again. I take it that the bulk of everyday manipulation is of the latter sort, as people keep playing.

>Greg's thesis is that there are ways of controlling the behavior of
>other people against their will or without their awareness other than
>by the use of overwhelming physical force.

Again, not "control" in the PCT sense, and, again, not "against their will."

>My thesis is that such apparent control (a) is an illusion, or (b) is of no
>importance to the "victim", or (c) is carried out only with the aid of the
>"victim," or (d) rests on a threat of force in the background. I contend that
>a properly functioning control system cannot be controlled from outside, so
>that any apparent success by a manipulator results from some defect of
>organization in the control system.

In sum, I am exploring (with your aid) the PCT foundations of (c). There is never any PCT-control "from the outside," is there? Success in manipulation depends on proper functioning of the influencee's control system (note that con men complain about the occasional mark "too stupid" to understand the game); it depends most crucially on the influencer having a reasonably good model of the influencee's control structure -- and this is one of the main topics for investigation in a GTM based on PCT. More precision in our definitions has great value, I think, for minimizing confusions due to this or that word meaning one thing to one person and another to someone else, so let's keep at it, OK? I'm learning a great deal from this.

Many happy returns (another manipulation?),

Greg

Date: Sun Aug 30, 1992 10:53 am PST
Subject: Misc as listed:

Subjects: Demos needed for PCT promotion
Musings on science
American Psychology
Greg: Prices on CSG literature also: What is Man?
Martin's paper

[From Dag Forssell (920827)

Recent comments by Avery and Rick:

>I think we will all continue to be frustrated on this subject until we

>come up with experiments that demonstrate this "absolute need."

How can control theorists be so off? We do not need more startling demonstrations. PCT tells us that all action is initiated by error signals.

What we need is to address the error signals that lurk out there in people. A synonym for error signal is dissatisfaction. We need to reach people who are dissatisfied with what they can accomplish, people with a yearning for something better. A better way to deal with each other.

A dissatisfied person will be open to suggestions and interested in trying a different solution.

Much of the debate on this net addresses people (directly and indirectly) who are perfectly satisfied with what they know, proud of it and ready to defend it.

Forget it. Ask people what problem they are anxious to solve. Ask if they are willing to think for themselves and evaluate an alternative. When people refer to authorities, they are not prepared to think for themselves. PCT does not need anything more than a student who is willing to think for him/her self and make the effort to understand the evidence.

Our challenge is to tell our story so that people become aware of the error signals they frequently deal with, and understand that we have a permanent solution they may like if they spend a little time looking at it.

About science:

PCT is a hard science. We expect 100% prediction and get 95%+, with the remaining 5% accounted for by expected imperfection of control - less than infinite loop gain and sloppy connections in the environment.

People schooled in the soft sciences have low or no expectations of prediction. Sometimes they do a poor job of describing what they are studying, much less offer explanations. I was astounded a while back in a personal conversation with a prestigious Russian Psychology PhD when he said matter of factly that the science of psychology has nothing to do with the practice of psychology. This is the way science is! I have since had others confirm this. To me, a science that has nothing to do with the reality of what it purports to study is no science at all.

People schooled in the hard sciences tend not to think of the softies as scientists at all, and the soft scientists don't begin to understand the distinction.

Yet all are "scientists" in the Kuhnian sense. Everyone observes the world through their own paradigms. No-one knows the Boss Reality. By Kuhn's definition as I understand it even a babbling child is a scientist. But there are differences in the standards the scientist tries to live up to. There are degrees of science rigor.

Modeling allows you to work to a high standard. It allows you to test your predictions and will prove you wrong in a hurry if you are off even a little.

Verbal exercises can be carried on indefinitely without any tests ever possible.

Last week I visited a friend and saw a few issues of American Psychology. A Swedish poetic quote comes to mind as I scan much of the content: "Up fly the words, thought is at a standstill". (Poetry lost here, sorry).

The following from the Comment section of the August 1992 issue may be of interest:

DISUNITY IN PSYCHOLOGY: CHAOS OR SPECIATION?

Richard J. McNally _Harvard University_

In his recent _American Psychologist_ article, Staats (September 1991) expressed concern about an increasing fragmentation in psychology that has produced a "crisis of disunity" (p. 899) exemplified by "great and increasing diversity--many unrelated methods, findings, problems, theoretical languages, schismatic issues, and philosophical positions" (p. 899). According to Staats, unless we unify the field, psychology is unlikely "to be considered to be a real science" (p. 910).

Although Staats cited Kuhn's (1962) early work on preparadigmatic science to support his thesis, Kuhn's (1991) recent work provides a more optimistic perspective on psychology's diversity. In a recent address based on his forthcoming book, Kuhn (1991) argued that cultural practices (e.g., religious, military, scientific) undergo a process akin to biological speciation. Following revolutions in science, new "species" emerge that develop their own research agendas, concepts, methodological standards, journals, and professional societies. Communication and cross fertilization remain possible when these subspecialities share intellectual ancestors but incommensurability arises as the tree of science branches outward, producing new limbs that share increasingly fewer roots. Although unity may occur within specialized domains of inquiry, the absence of an overarching framework has not impeded the progress of science.

To illustrate his point, Kuhn (1991) noted that physicists could once absorb new research by reading _Physical Review_ But today _Physical Review_ has fractured into four journals, and rare is the physicist who has the expertise, the time, or the interest to follow developments in more than one or two of these highly specialized outlets. Yet despite its fragmentation into subspecialities, physics has retained its progressive character. According to Kuhn (1991), mature science is a "ramshackle structure" whose semi-autonomous research communities develop theories that do not "sum to a unified picture of the world."

Kuhn's current views suggest that psychology's diversity may indicate vitality rather than impending demise. What Staats (1991) saw as a crisis of disunity may benignly reflect the natural history characteristic of cultural practices in general and science in particular. Moreover, developments applauded by Staats as exemplars of unification might best be construed as instances of further speciation (e.g., interfield theories; Bechtel, 1988). Fields such as biochemistry and cognitive neuroscience have not emerged through the unification of their parent disciplines; they have emerged through cross fertilization at the interface of neighboring disciplines. The result of such cross fertilization is not greater unification but rather greater specialization. Accordingly, biochemistry and cognitive neuroscience have developed their own respective research agendas journals, and professional societies. Finally, the Society for Studying Unity Issues in Psychology itself constitutes yet another example of speciation. Despite its goal of unifying psychology, this society exemplifies the unavoidable trend toward specialized inquiry.

In summary, Kuhn's (1991) view of science implies that diversity in psychology may signify vitality rather than centrifugal disintegration. Moreover, it may be neither possible nor necessary to unify all of psy- chology under the rubric of a general theoretical framework. Although efforts at unification ought not to be discouraged, the future of psychology is unlikely to depend on the success of such endeavors.

REFERENCES

Bechtel, W. (1988). _Philosophy of science An overview for cognitive science_. Hillsdale, NJ: Erlbaum.

Kuhn, T. S. (1962). _The structure of scientific revolutions_. Chicago: University of Chicago Press.

Kuhn, T S. (1991, November). _The problem with the historical philosophy of science_ (The Robert and Maurine Rothschild Distinguished Lecture). Address delivered in the History of Science Department, Harvard University, Cambridge, MA.

Staats, A. W. (1991). Unified positivism and unification psychology: Fad or new field? _American Psychologist_. 46. 899-912.

[Not understanding the difference between hard and soft science, the author does not recognize the difference between a) four branches of physics, each of which aims for 100% and holds itself to a standard of 99.99999% predictability, (which allows us to generate atomic power and send out a Mars lander), and b) the prattle of four or more branches of psychology, each of which holds itself to a 0-15% standard of predictability ("because that is the best we can do"). Psychologists talk about both Kuhn and Popper, but choose not to hear what they say.

The "crisis of disunity" does indeed signify impending demise.]

Another gem, this time from the May 1992 issue:

IS THE BEST ALWAYS PREFERRED?

Marilyn Freimuth
The Fielding Institute,
Santa Barbara, CA

On what grounds do we chose one theory over another? According to Howard's (March 1991) constructive realism, "The ultimate criteria for acceptance of one theory over others rests in each theory's ability to satisfy the set of epistemic criteria" (p. 188), which includes predictive accuracy, internal coherence, external consistency, fertility, and unifying power. To use Howard's metaphor, the theory that best follows the rules of scientific storytelling will be the theory we endorse.

As psychologists, we acknowledge the _conventions_ (i.e., epistemic criteria) on which one theory can be judged to tell a better story than another.(1) However, do these criteria become the basis for our theoretical preferences? In other words, Howard (1991) assumed that the "best" theory according to the rules of science will be the preferred theory.

Epistemic criteria seem relatively unimportant when graduate students in psychology select a theory. For a number of years, I have led a discussion in which doctoral students select from among 13 alternatives their most and least preferred explanation for a psychological event. Their choices are examined in terms of the criteria that Howard (1991) outlined. It is surprising how often the preferred explanation falls dramatically short on these criteria. Yet, NOT ONE of the more than 100 students who did the exercise has ever changed his or her position when presented with this information.

One could argue that psychology graduate students have not been fully so-cialized to recognize a good psychology story. But do they really act so differently from their teachers? What would lead a psychologist to prefer a new theoretical story? Howard (1991) referred to three reasons: (a) Research decreases a theory's predictive accuracy, (b) new theoretical developments decrease a given theory's external validity, and (c) a more "powerful" theory is developed that "tells a more compelling theoretical story" (p. 188). If in using the term _compelling_, Howard is referring to something other than satisfaction of epistemic criteria, he does not let on. Instead he goes on to assert that "whether or not a scientific theory initially feels right has _not_ become an important guide in theory choice" (p. 189). However, he does note that feeling right (e.g., empathic resonance) is an important rule for telling a good literary story.

In making this contrast, Howard (1991) missed a major implication of his own metaphor of psychological theory as story. As psychology, theories will be evaluated by the rules of good scientific storytelling (i.e., epistemic criteria). However, as stories, theories also will be evaluated as literary products, and as a result, nonepistemic criteria, such as "feeling right," will affect preferences. For example, in the same issue, Cushman (1991) argued that the appeal of Donald Stern's work does not reside with its being more "scientific" than other theories, but "his ideas feel _right_ [italics added] to many psychologists because they seem to capture the essence of their human experience" (p. 217). Other examples of nonepistemic factors influencing preference for a theory can be found in Gergen (1985), Prilleltensky (1989), Scarr (1985), and Harris (1979), who shows how the need to tell a good story about psychology's history may lead one to ignore a theory's failure to meet epistemic criteria.

To recognize that nonepistemic factors enter into theory choice does not mean a return to what Howard (1991) called "anything goes" relativism. Rather it behooves us to define the criteria that make a theory compelling to a person.

Most attention has been given to social and political factors. My own work (Freimuth, 1991) suggests that preferences are in part dependent upon a fit between a theory's basic premises and one's more general assumptions about how the world works. Other factors that could be studied include a fit between personality and theory (see Andrews, 1989, for a possible example), the role of a special teacher, or one's early experience or value system. This approach to choosing a theory is par- allel to the one Howard proposed for thinking about patient-therapist matching. Just as simplistic models cannot capture the complexities of the latter relationship, the matching of a psychologist with his or her choice of theory is multidetermined and not limited to how well a theory meets epistemic criteria. The previous points should not take away from how well Howard (1991) highlighted the implications of a narrative approach for thinking about different domains of psychological inquiry. However, as I have argued, Howard has not fully drawn out the implications of this perspective for how psychologists act when choosing their preferred theory or story.

Footnote: (1) It is not clear from his article whether Howard (1991) would agree that the rules of scientific storytelling are governed by some higher order story (i.e., epistemic criteria are relative and agreed upon conventions) or whether he sees these criteria as representing some necessary truth about the nature of science.

REFERENCES

Andrews, J. D. W. (1989). Integrating visions of reality: Interpersonal diagnosis and the existential vision. American Psychologist, 44, 803-8 1 7.

Cushman, P. (1991). Ideology obscured: Political uses of the self in Daniel Stern's infant. American Psychologist, 46, 206-219.

Freimuth, M. (1991). Pepper's world hypotheses and preference for psychological theories. Manuscript in preparation.

Gergen, K. (1985). The social constructionist movement in modern psychology. American Psychologist, 40, 266-275.

Harris, B. (1979). Whatever happened to Little Albert? American Psychologist, 34, 151-160.

Howard, G. S. (1991). Culture tales: A narrative approach to thinking, cross-cultural psychology, and psychotherapy. American Psychologist, 46, 187-197.

Prilleltensky, 1. (1989). Psychology and the status quo. American Psychologist, 44, 795-802.

Scarr, S. (1985). Constructing psychology: Making facts and fables for our times. American Psychologist, 40, 499-512.

[To say that "(a) Research decreases a theory's predictive accuracy", tells us that the term "theory" means something totally different to a soft scientist as compared to a hard scientist. Other commentary in the same issue seems to seriously suggest that theory and narrative are synonymous.]

In my career as a mechanical engineer, I have made literally hundreds of major predictions based on hard theory in the form of designs of tooling, parts and comprehensive products and manufacturing processes. (My major effort was the conception, design and build of a portable, inflatable, 14 foot catamaran sail and powerboat in collaboration with Hank Folson, whom I am proud to have sponsored onto this net). Each of these predictions has been tested by being built. Some have failed, some succeeded. Never once did it ever occur to me to resort to statistics to excuse a failure. Either a prediction works 100% or it does not. If it does not, you start over.

I am suggesting that those who have a soft science background and are wrestling with PCT may have a greater personal challenge than those with a hard science background. You will never arrive at an understanding from narrative, no matter how patiently put forth and repeated in many guises by Bill Powers, Rick Marken and others. You need to experience PCT, in hard experiments and/or in your own life.

PCT lives up to the standards of hard science. To understand what it is about, one needs to change the criterion for predictable success to 100%, and carefully review the published literature, starting with _Behavior: The Control of Perception_. The computer demonstrations are an excellent learning tool. Spend time with them! You will never understand PCT by trying to relate it to the endless prattle of contemporary psychology.

Long ago, I heard the saying: The responsibility for teaching belongs to the teacher. The responsibility for learning belongs to the learner.

There is a large body of teaching materials available and patient coaching on this net. The rethorical question I ask is this: Are you personally satisfied with what you know now and how it works for you? Do you really want to learn a better way? Are you willing to do whatever study and re-thinking it takes to really understand?

CSG literature: Attention: Greg.

I am happy to report that Christine and I now have our first customer scheduled! With two more close behind, hopefully.

As part of our program of teaching PCT to industry, we will bring a "literature table" and include a price list.

For the CSG literature, we would like to offer convenient delivery on some key items, at least. We already have samples of everything and a few LCS I and Intro to psych. Greg, would you please post the current prices and shipment costs for all items you offer:

For each title:

Single copies:		Volume orders:		
Mail order		Quantity	Price/ea or discount	Shipment for lot
	or single here:	1	?	?
Price ?		5	?	?

Shipment ?	10	?	?
	25	?	?
delivery time?	50	?	?
			delivery time?

I expect to give this info to companies and offer them to buy from me as a convenience.

Greg, on another subject: _What is man_? by Mark Twain is copyright 1906 by J. W. Bothwell. Since you are well versed in copyright issues, perhaps you can tell me: Can this be freely copied and distributed? Would you like a disc with the story? I am scanning it for myself before I have to return the book to the library.

Martin paper:

Thanks for your paper. We have just been back a few days. Wanted to get this commentary out of my system. Have to write short essay on PCT and performance reviews for prospective customer. Then I will read your paper and get you comments before Sept 4, as you asked directly.

Best to all, Dag

Date: Sun Aug 30, 1992 1:39 pm PST
From: Hortideas Publishing / MCI ID: 497-2767
TO: * Dag Forssell / MCI ID: 474-2580
Subject: Reply

From Greg Williams (920830 - direct)

Hi Dag,

I think that several of the issues you raise in your latest post are intriguing and important. Since I'm finishing up getting the arm paper shipped to SCIENCE and also have become embroiled (!?!?) with Bill regarding "manipulation," I'll wait until others chime in before saying my 1/50th of a dollar's worth. I predict plenty of responses (oops, I mean controlled behaviors).

>I am happy to report that Christine and I now have our first customer
>scheduled! With two more close behind, hopefully.

Great!

>As part of our program of teaching PCT to industry, we will bring a
>"literature table" and include a price list.

Great squared (at least)!!

>For the CSG literature, we would like to offer convenient delivery on
>some key items, at least. We already have samples of everything and a few
>LCS I and Intro to psych. Greg, would you please post the current prices
>and shipment costs for all items you offer:

For all four titles:

Single copies are full price, postpaid, five or more copies (can be assorted) are 20% off, but you must pay (UPS) shipping per invoice.

LCS [I] - $16.50
LCS II - $22.00
IMP - $25.00
MR - $18.00

All except MR by Rick are in stock and can be shipped within a week (next day if necessary; we usually go to UPS each Monday). I photocopy MR five at a time, so if I've just run out, there might be up to a week to 10-day delay until I get to Lexington to copy more. For special quick shipping, you'll be charged the exact extra charges.

>Greg, on another subject: _What is man_? by Mark Twain is copyright 1906
>by J. W. Bothwell. Since you are well versed in copyright issues, perhaps
>you can tell me: Can this be freely copied and distributed? Would you
>like a disc with the story? I am scanning it for myself before I have to
>return the book to the library.

Sure, I'd appreciate a (perhaps illegal!) disk copy. It would take some research to see whether Bothwell or somebody else still owns the copyright to it. Depends on whether they renewed it and if the holders are still alive. I really wouldn't know the most efficient way to begin the necessary research. One thing you could do: look to see if a RECENT compilation of Twain treats it as public domain (check out what the copyright page says) -- I'd be interested in knowing.

Best wishes for more customers,

Greg

Date: Sun Aug 30, 1992 2:56 pm PST
Subject: ; Gragmentation

[From Bill Powers (920830.1545)]

Greg Williams (920830) --

>I thought we'd already reached an agreement that one person cannot
>control (in the PCT sense) another. Are you using "control" more
>colloquially here?

No, technically. Control systems don't control their actions; they control their perceptions. By applying a disturbance to a controlled variable and thus disturbing, or attempting to disturb, the perception of it, an outside party can cause the action of the control system to vary. It will vary, predictably, so as to remain equal and opposite to the disturbance. That is why the disturbance has essentially no effect.

So the outside party can control the amount and direction of the action that the control system is producing.

Speaking of controlling "things" is ambiguous. Controlling "a person" means controlling the value of some variable associated with that person. Some of those variables (the amounts and kinds of actions) are controllable if the controlled variable is known. Others (perceptions, controlled variables) are not.

>... my definition of manipulation did not specify that the influencer's
>desired perceptions be hidden from the influencee. An example is the diet
>center counselor who manipulates the dieter so that both influencer and
>influencee see fewer pounds on the influencee.

If the manipulee is aware of the manipulator's goal, the manipulee can simply adopt it and control for it. In that case the manipulator has no control at all. Acceptance or rejection of the goal is entirely up to the so-called manipulee. The diet counsellor will have no effect on the client's weight if the client does not volutarily adopt the suggestions of the counsellor.

>... what I was trying to get at was manipulation as making the
>influencee control (to SATISFY his/her interests or wishes) in such a
>way as to ALSO satisfy certain interests/wishes of the influencer.

How can you MAKE the influencee control for anything? If controlling for something satisfies the influencee's own higher goals, that will be done. Otherwise it won't. The influencer can suggest possible goals, perhaps even goals the influencee is glad to be told about. But nothing at all will happen as a result unless the influencee translates the verbal suggestion or the seen example into a personal perceptual reference signal and activates it. Only a higher system in the influencee can do that.

>Some of the consequences could be unknown to or unpredicted by the
>influencee; other (or sometimes all) consequences (including the
>influencer's desired perceptions which are the "point" of the
>manipulation for the influencer) could be known to or predicted by
the >influencee. The latter possibility is in no way excluded by my
>definition of manipulation.

If the consequences are known to the influencee, what is the influencer's role? Why doesn't the influencee simply set that consequence as a goal and achieve it? What an influencer or manipulator wants is to be the AGENT that takes credit for the result. In order for that to happen, the influencee must NOT be the agent.

>Now, if you ask one of our kids whether Pat is providing them with "tasty"
>food so that she sees them eating "healthy" food, they'll say, "Sure."
>There's no deceit, indirection, misrepresentation of facts, false promises,
>or ambiguous communications here, but there IS manipulation as I defined it,
>because Pat altered some of the kids' perceptual variables in such a way as
>to result in them being able to control for certain perceptions which also
>result in Pat being able to control for certain perceptions.

This is completely analogous to the rubber-band demonstration of controlling behavior. By disturbing the knot, I can cause your finger to move to a preselected mark. You can be aware of this and it will make no difference, because you have decided that the most important thing to control is the position of the knot, not the position of your finger. Unless your kids have some reason to want Pat NOT to see them eating food she considers healthful, they don't care what she's controlling for by setting various kinds of food on the table as long as they like the result. Pat is not controlling them, or even disturbing or influencing them. Their controlled variables are undisturbed. They are the agents of their own actions. Whatever consequences of eating one kind of food rather than another Pat may imagine to be going on inside the kids, nothing is happening inside them that disturbs their goals one way or the other -- or their intrinsic variables either, one way or the other.

>Do you begin to grasp the pervasiveness of manipulation in everyday
>life?

I don't call this kind of interaction manipulation. It's hardly even interaction, in that the parties involved need do no negotiating and don't even have to be aware of each others' intentions. I think what is really pervasive in everyday life is a desire to have an effect on other people, and a willingness to believe that one has done so. In most cases this is quite harmless, as people can't affect each other (even each others' biochemical states) nearly as easily as they believe. When it becomes harmful, the manipulee generally knows it and resists.

In social interactions, the connotation of manipulation that is most common (to my knowledge) is that of control: one person determines what another person does. In my Random House dictionary, the first meaning is simply related to skill. The second meaning is "to manage or influence by artful skill: _to manipulate a person_."

>>Greg's thesis is that there are ways of controlling the behavior of
>>other people against their will or without their awareness other than
>>by the use of overwhelming physical force.

>Again, not "control" in the PCT sense, and, again, not "against their
>will."

What other kind of control is there but control in the PCT sense? If the manipulator wants to have a certain effect on the manipulee, and carries out actions that have that effect even despite disturbance, isn't that just plain garden-variety control? Would a manipulator be satisfied to go through the output motions without ever being able to see if the wanted effect occurred?

If the control is not against a person's will, then it is control of something that person considers unimportant, like the position of the finger in the rubber-band game. Of course at some point the position of the finger may become important to the manipulee, as when the finger approaches a hot spot. The control by the manipulator will then abruptly be lost.

>Many happy returns (another manipulation?)

Thanks -- that's what I wanted to hear.

Dag Forssell (920830) --

A most thoughtful essay. I find nothing to take issue with except the thesis that fragmentation is a good thing:

>In summary, Kuhn's (1991) view of science implies that diversity in
>psychology may signify vitality rather than centrifugal disintegration.

One good thing that comes out of fragmentation is that when everybody is wrong, fragmentation prevents one super-wrong idea from taking over. But science also requires at least some kind of "overarching framework" if it's to shift into a higher gear as physics and chemistry did.

A total lack of diversity implies loss of creativity. An overabundance of it indicates disorganization. Diversity is the natural outcome of reorganization; but reorganization is faulty when it goes so fast that the good results as well as the bad are lost in the next hasty reshuffling. The desired outcome, in any case, is not endless reorganization, but arrival at a better organization.

The paper by Marilyn Friemuth is a gem. Thanks for your great patience in keying (scanning?) in all this good stuff.

Altogether a welcome post.

Best to all, Bill P.

Date: Sun Aug 30, 1992 7:20 pm PST
Subject: More manipulation

From Greg Williams (920830 - 2)

>Bill Powers (920830.1545)

>Control systems don't control their actions; they control their perceptions.
>By applying a disturbance to a controlled variable and thus disturbing, or
>attempting to disturb, the perception of it, an outside party can cause the
>action of the control system to vary. It will vary, predictably, so as to
>remain equal and opposite to the disturbance. That is why the disturbance has
>essentially no effect. So the outside party can control the amount and
>direction of the action that the control system is producing.

More precisely, it seems to me, the outside party can control his/her perception of the amount and direction of the action produced by the control system. Right? (An aside: Of course, a disturbance can have TREMENDOUS effects on the control system's outputs -- but not (assuming good control) on the outcomes. I think it can be quite misleading to leave the word "effect" unqualified; in particular, it can lead one to think (falsely) that an alteration in an UNcontrolled variable has "essentially no effect" in general, when it in fact can result in switching of control to entirely different percepts than were being controlled before the alteration.)

>Controlling "a person" means controlling the value of some variable
>associated with that person. Some of those variables (the amounts and kinds
>of actions) are controllable if the controlled variable is known. Others

>(perceptions, controlled variables) are not.

OK, now I think I understand your definition of "controlling another person": it is controlling one's perception of certain actions produced by another person. Correct? If so, I see only one problem in replacing my "manipulation" with your "controlling another person." (See my next comment below.) My confusion about what you meant was due to the discussion on the net a few months ago wherein most everyone agreed that one cannot "causally" PCT-control others' controlled perceptions, as you explicitly note above, simply because the would-be controller does not have direct access to others' reference signals, which causally determine their (the others') controlled perceptions. If you look back at my "Prolegomenon," you'll see that I was trying to be careful to deny that one can alter the current reference-signal structure of another (but it is an altogether different question whether one can, by altering another's perceptions -- those which CAN be altered, of course -- play a role which contributes to a different FUTURE reference-signal structure than would otherwise exist; but that is at a tangent to my present concerns).

>If the manipulee is aware of the manipulator's goal, the manipulee can
>simply adopt it and control for it. In that case the manipulator has
>no control at all.

If the manipulee adopts and controls for the manipulator's goal, the manipulator doesn't necessarily stop controlling for that goal. Why should the manipulator stop? I suspect that usually BOTH manipulee and manipulator would continue controlling (for their respective perceptions of achieving the manipulator's goal). And the upshot, then, would be that the manipulator has an easy time of controlling, rather than no control at all. To say that the manipulator would have no control in this case seems to me to require an addition to your definition of "controlling another person," namely the requirement that such control cannot result from the controllee adopting the controller's goal overtly. I had no such requirement in my definition of manipulation, but I had another requirement, namely that the manipulator's altering some of the manipulee's perceptions helped to determine the manipulee's control processes which resulted in certain of the manipulator's perceptions matching their reference signals. So, in the case you describe above, given our respective definitions (yours not amended), it appears that the influencer (I can't say manipulator) has control but is not manipulating. So our definitions part company, right?

GW>>... what I was trying to get at was manipulation as making the
GW>>influencee control (to SATISFY his/her interests or wishes) in such a
GW>>way as to ALSO satisfy certain interests/wishes of the influencer.

BP>How can you MAKE the influencee control for anything?

Gloss it as "provide the opportunity for" control of perceptions which would not have been controlled for in the absence of the influencee's perceptions being altered by the influencer.

>If controlling for something satisfies the influencee's own higher goals,
>that will be done. Otherwise it won't. The influencer can suggest possible
>goals, perhaps even goals the influencee is glad to be told about. But
>nothing at all will happen as a result unless the influencee translates the
>verbal suggestion or the seen example into a personal perceptual reference

>signal and activates it. Only a higher system in the influencee can do that.

Yes, that's what I say, too. In manipulation, control of certain perceptions by the influencee results in control of certain perceptions by the influencer. The influencer wouldn't see what he/she wants unless the influencee were controlling certain perceptions (possibly with no knowledge of the influencer's desired perceptions). But for it to be manipulation, WHAT the influencee is controlling for must be partly the result of actions of the influencer. Now, of course, I don't mean to say that the influencer alters the reference-signal structure of the influencee, but rather that the influencer, by altering "don't-care" perceptions of the influencee, provides the opportunity for the influencee to control perceptions which the influencee would otherwise not have controlled. To repeat an example, "Look at this wallet I found here." No wallet, no control for taking the money in it. But the potential (so to speak) for control for taking the money was already there in that "good" mark who has shown "greediness" (not so subtly!) to the con man. So the con man has to work with what he finds in the mark's control structure, but he alters the mark's perceptions in ways so as to allow the expression of control by higher systems in the mark to benefit him, the con man.

>What an influencer or manipulator wants is to be the AGENT that takes credit
>for the result. In order for that to happen, the influencee must NOT be the
>agent.

By my definition, all the manipulator wants is for certain of his/her perceptions to match his/her reference signals. The perceptions need not include that he/she is "the AGENT that takes credit for the result." The influencer, by my definition, doesn't necessarily want ANYTHING -- influence can be "accidental." Con men know full well -- and say it -- that the mark gets the credit for his/her actions. Even the best con men repeat the old adage that "you can't cheat an honest man [or, presumably, woman]." (Actually, the old adage doesn't mean that one can't persuade an honest man to part with some of his resources by deceptive means; it means that it is best to con marks into acting dishonestly so they won't tend to run to the police when they realize -- if they ever do, some never do -- they've been "had.")

>Unless your kids have some reason to want Pat NOT to see them eating food she
>considers healthful, they don't care what she's controlling for by setting
>various kinds of food on the table as long as they like the result. Pat is
not
>controlling them, or even disturbing or influencing them. Their controlled
>variables are undisturbed.

By my definitions, she is CERTAINLY influencing them. Without Pat, there would be no tasty noodles. She alters their perceptions. Further, if she provided tasty "junk" (her definition) food instead of tasty "healthy" (her definition) food, they would gobble the "junk" food, and she would not be able to control her perception of seeing them eat "healthy" food. The situation certainly seems to meet the criteria of my definition of manipulation. The kids' controlled variables are undisturbed, yes, as I've been saying all along. WHICH percepts they are controlling (with their UNDISTURBED reference levels) are different because of Pat. And Pat's goals are met because (in part) of what Pat did; had she done otherwise (as suggested above), her goals would not have been met.

>They are the agents of their own actions. Whatever consequences of eating one
>kind of food rather than another Pat may imagine to be going on inside the
>kids, nothing is happening inside them that disturbs their goals one way or
>the other -- or their intrinsic variables either, one way or the other.

Yes, I agree. They are, in my terms, manipulated agents of their own actions.

>In my Random House dictionary, the first meaning [of "manipulation"] is
>simply related to skill. The second meaning is "to manage or influence by
>artful skill: _to manipulate a person_."

That certainly sounds like what I'm trying to investigate, using PCT ideas. No connotation of negativeness; skill is needed; influence is presumed -- it's all there, waiting for a theoretical underpinning.

>What other kind of control is there but control in the PCT sense? If
>the manipulator wants to have a certain effect on the manipulee, and
>carries out actions that have that effect even despite disturbance,
>isn't that just plain garden-variety control?

I think I'm straight on this now. Thanks.

>If the control is not against a person's will, then it is control of
>something that person considers unimportant, like the position of the finger
>in the rubber-band game. Of course at some point the position of the finger
>may become important to the manipulee, as when the finger approaches a hot
>spot. The control by the manipulator will then abruptly be lost.

With artful skill, some manipulators (or, if you don't like the word, simply "con men") can (almost literally) burn fingers without the manipulation ending. I'm still serious about understanding the nature and limits of this very real phenomenon, which I think is ubiquitous in family interactions, business dealings, education, therapy of all sorts, entertainment, religion, politics (robustly, as they say), etc., etc. I welcome additional suggestions. The literature on manipulation is surprisingly scanty, and mainly empirical.

Best, Greg

Date: Mon Aug 31, 1992 3:15 am PST
Subject: what is 'control'

[Hans Blom, 920831]

Although I have been listening in on this list for some years now, I have never actively contributed. I enjoy reading the list. You are a very creative and inspiring bunch of people. Being a control engineer, PCT is a 'so what' thing for me, nothing new. Some of the fields of application, psychology and (although implicit) philosophy, fascinate me. Allow me one contribution and I promise to keep silent again for a long time. Besides, it already takes me far too much time to monitor the list; I wonder how many hours there are in a day for some of the people who are the most active on this list. I admire you, but I cannot follow you: too many other things to do.

Some general remarks first. One: control is not everything. There is also a lot of non-control in the world. Two: what do we mean by 'control'? Do PCT and Skinner talk about different things or are they just different perspectives? Three: where does control come from? How does it originate? Four: when you are explicit and build models, the type of control that you use seems to be just the old-fashioned type PID-control. There is a lot in favor of PID-control, but a great variety of other types of control have been explored since: adaptive control, dual control, robust control, to name a few.

Also, when I read the things that you discuss in this list, I often notice that I see some things very differently and/or that I see different things. Take one of your popular examples: the control of movements by e. coli. I remodelled and reprogrammed this 'control system' from the descriptions that I found in the list discussions of the last few months. My model looks as follows: - The environment is a point source of nourishment (let's take sugar)

at x = 0, y = 0. The concentration profile in coli's (two-dimensional) environment is an inverse square law (different laws do not make much difference). Coli can sense the sugar concentration at the point where it is. The concentration at the point source is too high for coli (poisoning), far away it is too low (hunger). A radius of 10 is optimal for coli, i.e. coli will 'control for' a radius of 10. At a radius of 100 or more, coli cannot sense the sugar concentration anymore. Coli's initial position is at a radius of 50.

- Per iteration coli does the following:
  1. Coli tumbles, i.e. selects a new random direction.
  2. Coli swims, i.e. takes a step in the new random direction. The step size is inversely related to its sensed sugar concentration minus its optimum sugar concentration. This makes for large steps in case of hunger or poisoning, and small steps if coli feels fine. The step size is, however, limited to a maximum value (when coli is at a radius of 100 or more) and also to a (small) minimum value. This minimum value ensures that coli will always have to reestablish its position and cannot relax after reaching the goal. Alternatively, a small random displacement of its position (modeling physical effects of water flow and such) provides a similar kind of disturbance.

  Pick your parameters as you like (their values do not matter much), but make sure that coli needs at least some ten steps from a radius of 50 to a radius of 10 but not more than a few hundred.

Looking at coli's behavior in a number of simulation runs, I notice the following. In about half of the simulations, coli takes off into the blue beyond, where its sensors do not work anymore and where it is essentially lost, despite the fact that its initial position and at least its initial ten to twenty steps are within the area where its sensors do a perfect job. In the other half of the simulations, coli goes towards the radius of 10 and finds it or almost finds it. The path often looks very crooked with lots of moves towards, and then again away from, the optimal radius. Sometimes coli reaches a radius of 11 or 12 but subsequently moves away again and gets lost. In other cases, coli finds the radius of 10, lingers there for some time, escapes, returns, repeats this sequence a number of times, but eventually it escapes and gets lost in what for it is infinity. Looking at the simulation as it unfolds is an esthetically very satisfying experience, like art: you _almost_ believe that you 'understand'.

Now what I see is very reasonable behavior given the real environment in which a real coli lives. But it is not control. You might call it an attempt to control. Coli has more control than an inert protein molecule, but not much. Coli is not fully dependent on the Brownian movement of the molecules of the water it lives in: it looks as if it has some small say in the matter. I look at coli's behavior as a demonstration of _emergence_, of how control, in an as yet very primitive form, comes into being in the tiniest organisms.

It is not difficult to improve upon coli's control, but that requires additional equipment, either extra sensors or memory. Either provides a higher-dimensional view of the world and, given the right actuators, therefore also actions in more dimensions than before.

Better goal directed behavior results if I give my simulated coli a memory (just a single bit) that remembers whether the new concentration is 'better' than the old. If so, I make coli continue in the same old direction. If not, I make coli select a new random direction. The result is that the new coli (e. coli seems too small to have such a memory, but this could be the model of a larger sized bacterium) takes only very _short_ moves away from but _much longer_ moves towards its optimum environment. Eventually, this new organism reached the radius of 10 in all simulations, and each one stayed close to it forever.

But its 'towards' is not a 'steepest descent towards'. That is possible only with another sensor, with which a gradient can be established. When I give my simulated coli the capability to sense and use the sugar concentration gradient (again, e. coli cannot do that, but something the size of an amoeba can), it immediately takes the shortest path to its optimum and stays there ever after.

These simulations give a lot to think about. First, I see an _emergence_ of control, from no control at all (type 1; in a protein molecule, virus or small bacterium with no own modes of movement) to the primitive partial control that I see in my simulated coli (type 2) to the gradient descent (type 3) to the steepest descent (type 4). In reality, of course, there are no discrete types, and even an oxygen atom can be said to 'control for' the kind of chemical reactions that it will allow. This makes 'control' a rather fuzzy issue, very unlike a set of linear differential equations with fixed coefficients from which you can calculate P, I and D-terms.

Contrast penni sibun (920818.2000)'s general description of type 1-4 behavior

>i don't think behavior is ``what *needs* to be done.'' i think it's
>what *is* done.

with Rick Marken (920819.1000)'s more restricted description of type 4 behavior

>I have spent the last ten years trying to develop demonstrations that
>would show that precisely that assumption is wrong.

Is 'leaning on the world' a more appropriate term than 'control' for type 1 behavior? For which types of behavior is Rick Marken (920819.1000)'s

>The computer, using "the test for the controlled variable", can tell
>which of these behaviors is being done intentionally--so it is reading
>the subject's intention (mind)--hence, the program does "mind reading".

appropriate? Are often posed questions like "Why is there such reluctance on the part of those working on the hot approaches to behavior to even consider the possibility that behavior is the control of perception?" (this one from Rick Marken) showing a confusion about type 1 versus type 4 behavior?

Second, these four types of control have little to do with an 'increase of loop gain'. A larger loop gain for coli would mean larger strokes and the real danger of moving past a food source so fast that it cannot be located. Actually, coli's loop gain is very robust; a large range results in almost the same almost optimal behavior. Give it a too small loop gain and it cannot move; give it a too large loop gain and it gets lost immediately.

Third, the kind of control that is being discussed in this list is mainly the type 4 control. But even in humans I see all different types. Sometimes there is nothing you can do; you are just swept along with the winds of what is for you just random change. Sometimes you have a small say in the matter; you only have a vague 'holistic' sense the too complex situation and hardly know what to do (you cannot convert what you sense into meaningful actions). Sometimes you can tell when a new situation is 'better' but not the direction of 'best'. Sometimes you do know what is 'best', i.e. in which direction to go. All this has to do with the possibilities and limitations of your equipment, sensors, memory cells, actuators and the connections between them. In a sense, sensors and memory cells have the same function in that the latter can be viewed as sensors that provide access to (an encoding of) past experiences. Both increase the dimensionality of the 'impression' that an organism can have at any moment of time. If the dimensionality of the impression becomes too low relative to the complexity of the problem, when processing capabilities are too limited, when actions are futile, or when no clear goal exists, penni sibun (920825)'s

>why conduct elaborate deductions about yr surrounding when you can look and
>see? in particular, why maintain elaborate control structures when you can
>look and see what needs to be done? why make highly detailed Plans when you
>can improvise? why required instant expertise when you can improve by just
>keeping on doing it? why try figuring it out yourself when you can
collaborate
>w/ others who have been there? why insist on figureing out every situation
>afresh when you can trust yor accumulated experience?

strategy, sometimes called 'intuition', may work best. But note that the overwhelming majority of our experience has been accumulated through our evolutionary path through the eons, and that therefore our problem solving methods will be those that are best for _humans_ (in my opinion, even that is frequently too broad a generalization).

Fourth, we experience the behavior of simulated coli as extremely complex, despite the fact that the 'laws' on which that behavior is based are extremely simple. That is because coli's behavior is unpredictable, 'chaotic'. We cannot nicely formulate the link between the randomness of its steps and the emerged (limited) order except in subjective terms such as 'sometimes some colis seem to like being around the radius of 10'. 'Going down a level' for an explanation is, because it is so non-obvious, a never-ending investigation into the randomness of nature. Often, what remains are statistics, a subject that most of you seem not to like. But in physics, statistics is quite acceptable as a tool to obtain a sufficiently accurate picture of the world, such as in statistical mechanics which does not need quantum theory to arrive at 'emergent' quantities like temperature and pressure. Here, 'going up a level' is just as impossible, or simply too difficult for our limited human resources. The 'three body problem' from celestial mechanics is a classic example, and chaos researchers will tell you that 'almost all' problems are of this nature. It seems that the proportion of analyzable problems in nature is vanishingly small. We may have to 'kludge' forever.

When you only have a hammer, you see nails everywhere. I do not want to detract from the value of control theory (a hammer is, after all, sometimes a very useful instrument), but when you have control theory _only_, you see control systems everywhere. 'Control' is an elusive thing, just like 'temperature'. Control 'exists' because we want it to exist, maybe even because we _need_ it to exist to bring order into our chaos. In nature, the mechanisms that compose what we call a control system often seem to be thrown together from the lower-level stuff that happened to be available in an almost random, still badly understood process that we call evolution. Most emphatically, evolution has no goal; it just looks that way to the naive observer. If there is no goal, 'control' may be as naive a concept.

Avery.Andrews says

>This discussion is gotten completely out of hand, and, like Martin, I'm
>rather baffled by it. I suspect that people haven't done enough homework
>on the other guys' stuff to justify the things they're saying about it.

I agree. A lot of confusion and misunderstanding arises when people speak a different language. Learning each other's language is a prerequisite for an exchange of thoughts. Regrettably, this is a life- long process.

Hans Blom

Who am I? I studied electrical engineering, majoring in measurement and control theory. I work at a university and have done so for more than 20 years. My tasks are both teaching and doing research in 'medical engineering'. I have done work in a number of subjects: modelling, parameter estimation, control, adaptive control, dual control, man-machine interfacing (now called 'ergonomics'), usually applied to anesthesia and intensive care monitoring and control systems. My long-term project is called 'servo-anesthesia', but that has proven to be a very elusive goal. I have looked around in artificial intelligence, mainly expert systems, neural networks and genetic algorithms. Also psychology, physiology, and some philosophy. My Ph.D. work was the design of a real time expert systems toolbox (special purpose programming language, its compiler and its inference engine). An application was the design of an expert system based blood pressure control system, where the computer controls the flow of an infusion pump that delivers a drug in such a way as to stabilize the patient's blood pressure at a lower than normal level. The core of the system is a PID-controller, but that core is surrounded by a 'safety shell' that contains expert knowledge, both medical (what to do on special occasions such as shock) and about how to tune a PID- controller. Tuning must keep the controller both fast and stable in the face of large variations in the patient parameters. Difficult to handle is the variability of the patient's sensitivity to the drug, which can vary unpredictably by a factor of about 80. Most difficult is, however, that we cannot explicitly specify what is best for the patient.

What impressed me most in my research was the influence of the _unknown_. Some examples: how to learn the initially unknown characteristics of the patient as fast as possible (time, especially medical time, is money) when the controller is started, while guaranteeing safety (predictable, 'almost' optimal behavior); how to keep track of changes of the patient's characteristics while controlling stably (no test signals); and how to handle missing feedback measurements (a temporarily broken feedback loop) while still ensuring safety for the patient. That cannot always be done. Sometimes the system has to alarm and transfer control back to the physician. You have to design in what to do when more and more sensors fail. That is not easy.

My expert system's logic has three truth values: true, false and unknown. Ensuring that the system reaches a conclusion of either true or false ('is the patient OK, yes or no?') but not the uninformative 'unknown' when given the largest number of input 'unknown's generally does not seem a mathematically tractable problem. Still, we want to design solutions. Creativity seems the only method, but one that cannot be formalized, regrettably.

Again, control is not everything, and PCT is even more limited. Just one example. 'Dual control' theory considers optimization of control over some time in an uncertain world. It recognizes that, in general, control will be better the more accurate a model of the world is available. Thus, in many cases _active_ learning (system identification) is called for. Active learning, however, means introducing 'test signals' and thus a disturbance of the observation, which of course _degrades_ the quality of control. But only in the short run. Over a long time, the observations of the response of the test signals yield an improved model which provides a better control. PCT does not consider this active learning. Its 'reorganization' is, if I understand correctly, a random process that only occurs when errors remain large for some time. There is no provision to temporarily _create_ small errors in order to prevent later larger errors. This type of active learning ('curiosity', 'exploration', 'play') is pervasive in (higher) animals and humans.

Wittgenstein once remarked that people cannot think unlogically. For me this means that every remark, however stupid it initially seems, has its point and can be learned from (maybe that's why I read so slowly). Sometimes others have just a different perspective on the same 'truth'. Sometimes someone has an impression that we can show to be inconsistent with what science shows us. But even then I think that we have to take that impression seriously and try to find out where it comes from and what are the sources of misunderstanding.

Date: Mon Aug 31, 1992 6:51 am PST
From: mmt
MBX: mmt@ben.dciem.dnd.ca
TO: * Dag Forssell / MCI ID: 474-2580
Subject: Re: Misc as listed:

Dag,

>A Swedish poetic quote comes to mind as I scan much of the content:
>"Up fly the words, thought is at a standstill". (Poetry lost here, sorry).

"My words fly up, my thought remain below" (or something like that). The King, spied on by Hamlet at his prayers, just after Hamlet decided not to kill him because he was at prayer.

Same?

Martin

PS. Did you get my Paris abstract yet?

Date: Mon Aug 31, 1992 7:28 am PST
From: Dag Forssell / MCI ID: 474-2580
TO: Martin (Ems)
MBX: mmt@ben.dciem.dnd.ca
Subject: Reply

From Dag

Yes, I got your paper. Meant to indicate that in my post yesterday.

With this mornings posting, I have cleared my major error signals and will read your paper and comment by tomorrow at the latest.

No, I think the Swedish quote is an original. Close, though.

Regards, Dag

Date: Mon Aug 31, 1992 7:04 am PST
Subject: Sloman: The Mind as CS

[From: Bruce Nevin (Mon 920831 10:41:17)]

Interesting title of talk by Sloman, highlighted in the following:

| From: neuron-request@CATTELL.PSYCH.UPENN.EDU ("Neuron-Digest Moderator")
| Newsgroups: comp.ai.neural-nets
| Subject: Neuron Digest V9 #41 (conferences)
| Date: 27 Aug 92 04:53:43 GMT
| Reply-To: "Neuron-Request" <neuron-request@cattell.psych.upenn.edu>

| There follows the provisional programme of papers for the Royal Institute
| of Philosophy Conference "Philosophy and the Cognitive Sciences" at the
| The University of Birmingham September 11-14 1992.

| [I have clipped the left column containing dates and times -- BN]

| `Blindsight',

| Michael Tye

| `Wittgenstein and Connectionism: a significant
| complemenarity',
| Stephen Mills

| `Levels of description in connectionism: cognitive
| transition, dynamical system, neural implementation',
| Terry Horgan (and John Tienson)

| `On the notions of specification and implementation',
| Anthony Galton

| `The connectionism/classicism battle to win souls',
| Brian McLaughlin

|
/***
| *
*
| * `The mind as a control system', < I wonder what he
*
| * Aaron Sloman < means by this? -- BN
*
| *
*
|
***/

| `Systematicity in the vision to language chain',
| Niels Ole Bernsen

| `Two notions of implicit rules',
| Martin Davies

| `Do your concepts develop?'
| Andrew Woodfield

| `Quine, Simon and the pragmatist's project',
| Stephen Stich

| The total for Registration, Bed and Breakfast and all meals is 122.50
| pounds sterling. To obtain a registration form, please email
| rip92@bham.ac.uk.

Date: Mon Aug 31, 1992 7:11 am PST
Subject: Commercial error signal?

[From Dag Forssell (920830-2)
((920827) was begun 8/27 but posted earlier today).

A prospective customer asked for an explanation of our view of
Performance Reviews.

Any comments? Gets mailed Monday pm.

August 30, 1992

Dear Mr. :

It was a pleasure to meet you last Monday. Our flight to Los Angeles was pleasantly uneventful, and we enjoyed reading the Inspiring Quality Employee Guide. Your receptionist also gave us the 1991 Annual Report.

The past week has been hectic. I am pleased to report that we now have received our first booking for a complete presentation of our programs.

You asked us to provide a proposal on our alternative to the employee development / performance review program.

Our approach to performance reviews is built into our programs as one of the "Leadership applications" on the second day. I will attempt here to provide a thumbnail sketch of our basic approach and then show the carryover to Performance Reviews.

You will recognize the following diagram from page two of the separate summary: Control: What it is; where it applies:

(Control Diagram)

Taken as a simplistic summary of a person, this diagram shows that what a person wants is of major importance, along with how a person perceives things. The arrows represent signals and influences between the elements labeled in this diagram. The arrow going from the comparison between Want and Perception is called "error signal" in engineering language (error = difference). This is a measure of satisfaction, where error = zero is satisfaction, and anything else a degree of dissatisfaction.

You have heard the old saying: "Necessity is the mother of invention". It is more accurate to say that: "Dissatisfaction is the mother of invention".

This measure of satisfaction is central to control. It drives everything, large and small. Hunger and thirst are physical examples of this "error" signal. They drive us to action, don't they?

Respect for a person, as we teach it, means to encourage and actively support the exercise of well-informed, effective and satisfying individual control. This applies to your family, associates, employees, vendors, customers and friends.

A challenge for management is to direct and coordinate the activities of many employees while allowing them to direct themselves. When you accomplish this, the company can indeed run itself, as you suggested you want your company to do.

From a PCT perspective, the name of the game is associate satisfaction at all levels in the organization. Satisfaction by definition is present when there is no difference (or only a small difference) between what the associate wants and what the associate perceives.

What then does the associate want, and what does the associate perceive? Both of these are subjective, unique to each individual, and subject to re-consideration by the associate. They are far from simplistic, but can be understood. You will find interesting parallels between your "Corporate goals waterfalling" and the PCT model's hierarchical structure in this regard.

If the associate wants the customer to be satisfied, and perceives the customer to be unhappy, then the associate will not be satisfied but will take action. If the associate wants the shareholder........ Associate satisfaction is all there is. Associates is all you have to work with. Associates must want and perceive many things and take effective action for the company to function well.

Our application in the first day program "Management Application" is called "Teaching Effectiveness." The thrust of our approach is to have the manager/coach (given a conflict situation) assist the associate in a re-consideration of wants and perceptions so that the associate is capable of functioning effectively = with satisfaction.

This is accomplished by asking the associate to spell out:

1) what the associate wants and

2) what the associate thinks the company wants.

The manager can gently question the wisdom of some of the personal wants and clarify the company's wants.

3) The associate must want to resolve any incompatibility, but again can get support and counsel on how to do that from the more experienced manager.

4) The associate develops a plan with the manager's help and

5) both follow up.

To reiterate: The associates are in complete control (self-direction) and are given every opportunity to align in their mind their wants with the company's wants.

By respecting the individual associate's self-direction so completely, and coaching the associate openly in the associate's own best interest, trust is developed. This results in personal and professional growth. Traditional performance reviews have the best of intentions. They are designed to provide feedback and help associates improve and develop. That is like motherhood and apple pie. You cannot argue with it. Everyone will agree that performance reviews are desirable. But everyone hates them. Why?

Please look again at the chart. At the bottom is an element in the environment labeled: Disturbance. Disturbances affect the things (Variable) we care about (Perceive, pay attention to) all the time. We counteract their effect on the variable. A disturbance so strong that we cannot overcome it overwhelms our control capability. We are unable to counteract it effectively. This is violence. It is not satisfying. We cannot eliminate the error signal. A chronic error signal is chronic stress.

When an associate or (as planned) supervisor is presented with a judgement as a fait accompli, the person is presented with a disturbance (of some aspect of self-concept) that cannot be countered in any effective way. At the same time the judgement does not give you meaningful detail on how to improve. It does little to enhance the person's capability to perceive or choose wants effectively. It does little to make the person more effective and capable of satisfying him/her self. It will not in any way create trust to have the judgement be an anonymous, supposedly objective compilation of many judgements. That will just make it that much more impossible to counter.

An approach to performance reviews follows naturally as an extension of the "Teaching Effectiveness" methodology.

While I have told you that the Purposeful LeadershipTM programs are new, (my contribution is in the clarity of illustrations and a unified methodology), the approach has been tried before and has worked well wherever it has been implemented. It accomplishes the following:

1) On the job satisfaction for associates

2) Supervisors get a thorough and realistic picture of what the associate is doing, is capable of doing, and where assistance is needed.

3) It teaches the associate how to perform better. - To be more satisfied.

This requires an interactive process and must of course be respectful, leaving the associate in control.

This process should be as continuous as possible. This review should ideally be conducted once a week, at least once a month.

The suggested process goes as follows:

Schedule a regular, undisturbed meeting.

1) The associate initiates the review by submitting a (hand) written description of one or two projects, challenges, situations s/he has dealt with since the last meeting. S/he can orally embellish any details.

This leads to a supportive and appreciative discussion, focusing on current job issues of whatever kind.

2) The associate describes one or two areas where s/he believes improvement would help improve performance.

3) Committment to resolve is not an issue where there is no conflict.

4) The supervisor works with the associate to formulate a plan for improvement.

5) The associate and supervisor both commit to carry out the plan.

When reviews are perceived as normal, non-threatening, commonplace events, people will relax. They will talk freely about their expectations, goals and disturbances (problems). The focus is on assisting one another to be more effective and satisfied. This is a very personal interaction with mutual respect. This is not institutionalized. This does develop trust.

You will note that when it comes time to consider a promotion, there is an excellent and very personal record of the associate's capabilities and progress.

In our meeting, you noted that the basic values we have expressed seem in alignment with yours. I can certainly see why when I read your Corporate Statement of Values.

I really appreciate the detail and spirit expressed here.

We will be pleased to present an overview of our program in four hours as you mentioned. If you allow us to present the whole program as a pilot presentation to a representative group from the Quality Council and/or Team-on-teams, some lively review of the employee guide will result as they consider their new insight and relate it to the programs.

Benjamin Franklin is quoted as saying: "People are best convinced by that which they themselves discover." Your quality group will stay satisfied and in control while they voluntarily incorporate new information that allows them to perceive better and select wants better in their effort to improve the Inspiring Quality Employee Guide.

Thank you for this opportunity to explain a bit more of how the PCT paradigm sheds light on the thorny issue of performance reviews.

Please let us know if this commentary meets your Valid Requirement and ask us for any other supportive information you may require.

Sincerely,

Dag Forssell

Date: Mon Aug 31, 1992 7:38 am PST
From: mmt
MBX: mmt@ben.dciem.dnd.ca
TO: * Dag Forssell / MCI ID: 474-2580
Subject: Re: Reply

I sent my note before reading all the way through your posting. Yes, you did indicate at the end that you had received my Paris abstract. I shouldn't shoot from the lip so often.

Martin

Date: Mon Aug 31, 1992 10:14 am PST

Subject: Marken Computer Demos/Control

[From Rick Marken (920831.1100)]

I've been off the net for the weekend. Amazing how much work one can get done when you don't try to respond to every disturbance created by net postings.

Just a couple quick comments:

Gary Cziko (920828.2237)

>Rick, some time ago you mentioned the possibility of "packaging" your
>programs into one with a menu interface. I would find this very useful for
>my students.

Well, then I will make it a shorter term goal to do this. It probably would fit in with whatever I do for Dennis.

Hans Blom (920831) --

>Some general remarks first. One: control is not everything.

Right. "The test for the controlled variable" is how you determine whether or not control is involved in any observed behavior.

> Two: what do we mean by 'control'?

Production of consistent results in the face of disturbnce.

> Three: where does control come from?

Control systems.

>How does it originate?

God.

Actually, I don't understand what you mean.

>Four: when you are explicit and build models, the type of control
>that you use seems to be just the old-fashioned type PID-control.

We build models to accurately mimic behavior, not to keep up with scientific or engineering trends -- although we like to keep abreast of the latest technologies since they might make it possible to improve our models (in terms of their ability to mimic behavior).

By the way. Your description of your e. coli model was somewhat puzzling. When, for example, does your e. coli model change to a new direction? What is the criterion for change? In our models, a "counter" is set after each tumble. The rate at which it counts down depends on the currently experienced gradient of attractant. The model works beautifully and it does, unquestionably, control it's perception of gradient. See my paper (with Powers) in my "Mind Readings" book (it's the one on Random -walk Chemotaxis).

>When you only have a hammer, you see nails everywhere. I do not want
>to detract from the value of control theory (a hammer is, after all,
>sometimes a very useful instrument), but when you have control theory
>_only_, you see control systems everywhere.

That would be true if we did not have a formal criterion for determining whether or not we are dealing with a nail. "The test" is that crierion. The fact is, what people call "behaviors" are consistently produced results of the actions of organisms -- it is this consistency that makes it possible to name these results -- "drinking", "driving", "speaking", "walking", etc. Each of these words describes experiences in which some variable is brought to some predetermined state and maintained in that state against disturbance. That is, these are words that describe controlled variables -- the position of the drinking glass, the position of the car relative to it's lane, the grammatical order of the words in a sentence, the direction of one's body as the legs are moved. The variables controlled are not always obvious (because they are perceptual variables) but the fact that a consistent result is produced in the context of variable disturbances (and that variable actions are used to produce this consistency) suggests that SOME variable is being controlled.

> 'Control' is an elusive thing, just like 'temperature'.

I don't understand this. I can demonstrate control as easily as I can measure temperature. What's the problem?

> Control 'exists' because we want it to exist

No. It exists because that is the way systems work that exist in a closed negative feedback loop; control happens.

> Most emphatically,
>evolution has no goal; it just looks that way to the naive observer.

Evolution may have no goal, but organisms (which evolve) most emphatically do have goals -- which is the reason why they evolve.

>If there is no goal, 'control' may be as naive a concept.

If there is no goal, then there is NO control.

But organisms have goals (lots of them). They are goals regarding what they should PERCEIVE and these goals can be revealed using the test for the controlled variable.

Hasta luego Rick

Date: Mon Aug 31, 1992 10:43 am PST
Subject: Re: Manipulation

[Martin Taylor 920831 1330]
(Bill Powers 920829.2130) and (Greg Williams 920830)

I hope this is short. Time presses.

Bill. I think this is another occasion in which the language we use points in opposite directions while the thoughts are close. But I'm not sure.

>Martin Taylor (920829.1315) --
>
>>One should never say that one "controls" the location of the car
>within >its lane on the road. One should say that one controls the
>percept of >the location of the car....
>
>The stages of Satori. Before I understood, I said that I controlled
>the car within its lane. While I was understanding, I said that I
>controlled the percept of the car within its lane. After I understood,
>I said that I controlled the car within its lane (for I know only the
>percept in any case).

Well, yes. That's the normal way of looking at it. But when we are trying to be precise, and communicate exactly, then we have to go back to separating the controlled percept from the CEV, which is what I was trying to do. We do, easily but sloppily, use third-stage language most of the time.

>>We need a different word for what happens to the CEV corresponding to
>a >controlled percept. For lack of a better word, I propose here to
>use >X-control (externalized control). One can X-control passive CEV's
>in >the real outer world. One cannot control (P-control?) them.
>
>I resist this terminology. If the (hypothetical) external CEV (complex
>environmental variable) does not behave as the percept does, then the
>CEV has been misdefined; it must entail aspects that are not matched
>in the perceptual function of the controlling system.

As you put it here, ALL CEVs are and always will remain misdefined, inasmuch as there probably exists no construct whatever in the outer world that behaves precisely as the perceptual input function that defines the CEV does. All we can know is those aspects of the outer world that correspond to some perceptual
function, but our actions affect the real world as it is, not as we perceive it.
We perceive the effects our actions have on the real world as filtered through our perceptual input functions, and that filtering is both noisy (intrinsically)
and subject to error (contextual effects). The CEV we think we are controlling
corresponds only by fluke with any CEV controlled by anyone else, and even then we cannot know that this correspondence has occurred.

What we can know is that we can perceive someone else acting on the CEV we are trying to X-xontrol--i.e. we perceive that our percept corresponding to the
CEV is changing in coordination with some perception of action by someone else.
We might even, applying the Test, perceive that the other is X-controlling a CEV highly correlated with one we are controlling (or disturbing). We cannot

perceive or determine by any finite procedure what the other is P-controlling.

>I want to minimize
>reification of that hypothetical reality. To speak of a CEV apart from
>the perceptual system that defines it for a controlling system is to
>give the CEV an independent existence of its own, to objectify it, and
>to imply a third-party observer who can know the true form of the CEV
>which may be partly misrepresented in the controlling system. There
>may be many variables in the external world. But they are not
>organized into CEVs.

It is this paragraph that makes me think we are trying to say the same thing.
It is what I was trying to get across in making the P-control versus X-control
distinction. X-control is X both for X-ternal and for "hypothesized
variable."
X-control is based on the notion that what you perceive has some identity in
the outer world. Maybe it does, and maybe it doesn't. You can't know.

>There is only one form of control: control of the percept.

Yes, that's my main point. And using that as the sense of "control", we CAN
control other people, as you demonstrated in a posting last week.

==========

>A disturbance that has a transient effect only will not materially
>affect the behavior (output) of the control system. If the control
>system could react within the interval we classify as "transient," the
>disturbance would be successfully resisted and there would be no
>transient effect on the perception. Remember that the main means of
>producing predictable outputs from a control system is to apply a
>disturbance THAT IS CANCELED BY THE ENSUING ACTION.

You can't have it both ways. Either there is an error signal that results
from the disturbance and generates an output signal that resists the
disturbance, or there isn't. If there is, it is an indication that the
disturbance did affect the percept. If there isn't, we don't know whether
any percept was affected, but if one was, it remains altered.

>Control systems evolve and reorganize until they can react adequately
>on the time-scale typical of normal disturbances. Don't think of
>disturbances as transient "events." They are simply independent
>variables that influence the controlled variable. They may be brief,
>or like gravity they may be present permanently.

All feedback loops have transport delay. All linear feedback systems can
be characterized by their bandwidth and phase response. Non-linear ones
are harder to characterize, but among their characteristics are information
capacity measures. All of these apply to the dynamics. I think the word
"adequately" in your first line is key. It means that the organism will
survive, behaving this well, and that no excess resources will be applied
to ensure better behaviour than that.

>I repeat that an influence (or disturbance) is neither control nor
>determination. It is simply an influence, which may or may not have
>any effect, depending on other influences that may be present at the

>same time (such as the output of a control system).

Well, I refer again to your clarifying posting, which I had better look up...
Yep. Here it is (920828.0900)

>People are not generally content merely to apply an influence to
>someone else's behavior and accept the result. To do that would be to
>accept the fact that the other's behavior is most likely to proceed as
>before with no effect from the supposed "influence." What happens in
>reality is that if the first attempt at influence fails, as it is most
>likely to do, there will be continued attempts involving varieties of
>influences and increasing force behind the influences. The influencer
>reveals the fact that this is an attempt at control, not merely at
>influence. The object is to have a particular effect on the other's
>behavior that matches the influencer's goal for the effect.
>
>The transition from mild and innocuous attempts at influence to
>concerted attempts to control is inevitable, for the simple reason
>that mere influence has almost no effect.

The same applies when we try to "influence" a ball to go up an incline by blowing at it. If we just blow, it may go up and come down again, or it may go sideways, or something else. If we really want it to go up, we modify the direction and strength of blow according to our perception of what the ball is doing. No difference when we "control" another person. We keep trying different "influences" until the person does what we want or until we give up.

The original (confusing) claim against controlling another person was that we could not, from outside, reliably set a reference level for any ECS within the other person. This view is the one I was seeing as irrelevant, since we control only our own percepts. Even internally, no ECS controls another. An ECS controls only its own percept. It can not set the reference level for another ECS, though it can influence that reference level in conjunction with influences from possibly many other ECSs at its own level. And even if an ECS was the only influence on the reference level of some lower-level ECS, it would still not be controlling that ECS. If it could perceive directly the reference level that it was setting, it could (X-)control that, but the lower ECS might react in any number of ways, depending on the correctness of its organization, its gain, the environment within which it was trying to control, and so forth.

No, control of another has nothing whatever to do with setting reference levels
for another. Control of another may be achieved easily if it so happens that the other's reference levels come to lead to actions that are perceived as appropriate to the reference levels in the "controller." But the other's reference levels are a red herring in the discussion of whether one person can control another.

===============
Greg,

>>What is meant is one person controlling another person's actions
>>surreptitiously, leading to that other person doing something one could not
>>get the person to do simply by asking.

>
>I thought we'd already reached an agreement that one person cannot control
(in
>the PCT sense) another.

Well, I haven't reached that agreement (see above).

>In sum, I am exploring (with your aid) the PCT foundations of (c). There is
>never any PCT-control "from the outside," is there? Success in manipulation
>depends on proper functioning of the influencee's control system (note that
>con men complain about the occasional mark "too stupid" to understand the
>game); it depends most crucially on the influencer having a reasonably good
>model of the influencee's control structure -- and this is one of the main
>topics for investigation in a GTM based on PCT. More precision in our
>definitions has great value, I think, for minimizing confusions due to this
>or that word meaning one thing to one person and another to someone else, so
>let's keep at it, OK? I'm learning a great deal from this.

I quite agree, and that is why I introduced X-control to contrast with real (P-)control. And with P=control, one person most certainly CAN control another, whether it be by force (i.e. in a conflict situation) or by taking advantage of presumptions about the other's reference levels to set up disturbances that lead to the desired actions. Each method fits exactly within the normal PCT definition of control.

Martin

Date: Mon Aug 31, 1992 1:32 pm PST
Subject: e coli/anniversary

[From Rick Marken (920831.1430)]

By the way, I didn't finish my description of the e coli model.

I said:

>In our models, a "counter" is set
>after each tumble. The rate at which it counts down depends on the currently
>experienced gradient of attractant.

I should have added:

e coli continues to move (at a fixed rate) in its current direction until the counter reaches zero at which point e coli tumbles (selects a new direction
at random); the counter is reset and e coli proceeds in the new direction until the counter reaches zero again. Since the count down is faster when e coli is moving away from the target, e coli tends to spend more time moving toward the target and less time moving away -- even though it always moves at the same rate. The behavior of the model matches the actual behavior of e coli (as reported in experiments done by Koshland and others). The controlled variable is the gradient of attractant. The refernce for this variable can be anything from 1 (maximum gradient -- ie. attraction) to -1 (avoidance).

I was reading a biography of John B Watson this weekend -- called "Mechanical Man" by K. Buckley (Guilford, 1987) Very interesting look at the early politics of psychology. I noticed that 1993 will be the 80th anniversary of the publication of Watson's classic "Psychology as the Behaviorist Views it" which appeared in Psych Review in 1913. I think someone might consider submitting an article called "Psychology as the Control Theorist Views It" to Psych Review for publication in 1993.

Hasta luego Rick

Date: Mon Aug 31, 1992 1:45 pm PST
Subject: Exploration and Reorganization

[Martin Taylor 920831 17:20]

"All behaviour is the control of perception." But is it? I've been pondering a bit on Bill's comment a while ago that he found in his reorganization simulations that the flip of a sign in an output-to-reference connection caused a bad transient that troubled the whole net. It leads to the question "What is exploration?"

Exploratory behaviour is "what if" behaviour. One is not controlling for any particular reference (at a higher level one is controlling for a percept of knowing more than one does, but that's outside this point). One is just acting,
more or less randomly at the level at which exploration is being done. Of course it isn't random at lower levels, and probably not at higher ones.

Exploration could be considered action without a reference, or action with a severed connection between comparator and output, with a noise source substituted for the error signal. "What happens if I give a weak positive output?...OK, now what happens if I give a strong negative output?...OK, now what..." and so forth. The results could be used in for reorganization, in this case the adjustment of the signs of the output gain function or the signs of randomly selected output links. Flipping signs has no effect on the behaviour of the hierarchy if the signs are flipped in a loop with zero gain.

Exploration is for learning, and that is not usually coupled with a desire for the results to have particular values (unless you are a politician); so the actions involved in exploration are carried out in the absence of (strong) negative feedback. This makes reorganization safe, at the cost of providing unpredictable (but compensable) disturbances to parts of the hierarchy that are controlling.

In the Layered Protocol formulation, before I heard of PCT, I argued that the purpose of casual conversation was for the partners to develop models of each other that would permit them to communicate effectively when the need arose. In the same vein, I suggest that the purpose of casual exploration is to develop organizations that permit effective control when the occasion arises. Exploratory behaviour, then, is NOT the control of perception, but the discovery of how to control perception.

Martin

Date: Mon Aug 31, 1992 4:12 pm PST
Subject: anesthesia system

[Avery Andrews 920901.0957]
(Hans Blom, 920831)

Thanks for the very interesting posting. As a practicing cog. sci. softy (Chomskyan generative grammar), I always am interested to see how actual engineers look at things - there's something to be said for the idea that with the rise of cognitive science in the fifties, the engineers disappeared into the basement with the grownups' toolbox, leaving everyone else to play with the duplo bricks.

The anesthesia monitoring systems sounds like an instructive integration of thought and action - is there a comprehensive writeup anywhere?

Avery.Andrews@anu.edu.au

Date: Mon Aug 31, 1992 6:13 pm PST
Subject: Re: Agre's Cooking problem

[From Jeff Hunter (920831)]

[Re: Bill Powers (920827.0800)]

I'm planning to get around to a lot of posts, but I Have to comment on Agre's program (called Toast)

It was tasked with making a breakfast, and did so (in simulation) expeditiously, whereas a STRIPS planner blew up after 6 hours.

This was (if you pardon the pun) a cooked example. None of the subtasks interfered with any other in any major way.

To show why, here is an example of weekend goals:

- wear sunhat
- get new car tires
- spend weekend at cottage

And these are fed into Agre's weekend planner (main mechanism recapped below):

>While there are unsatisfied goals,
>   Choose aribtrarily an unsatisfied goal.
     ... use available tools and resources to satisfy goal.

[pick a goal: spend weekend at cottage]

- walk to car
- enter car

```
- start driving to cottage
        [slack time. try other goals]
        [cannot wear hat. cannot walk to closet. am driving]
        [cannot drive to repair shop. am driving to cottage]
- arrive at cottage     [goal #3 in progress]
        [pick a goal: wear sun hat]
- walk to closet        [at home!]
- wear hat              [goal #2 satisfied]
        [pick a goal: get tires changed]
- walk to car           [at cottage!]
- drive to repair shop
- exit car. wait for tire change
        [slack time. try other goals]
        [pick a goal: spend weekend at cottage]
        [car unavailable. no problem]
- walk to cottage
- arrive at cottage     [goal #3 in progress]
        [pick a goal: get tires changed]
- walk to car
- remove car from repair shop
        [pick a goal: spend weekend at cottage]
- drive to cottage
- enjoy weekend         [soaking your feet]
```

Since all goals require movement they all conflict. A bad choice of goal ordering can make for a considerable amount of walking!

Any decent planner can find the best ordering for these weekend goals in nearly the same time that Agre's one would find a lousy one.

The "planning weenies" have considered a lot of cases of conflicting tasks, and their work should not be discarded lightly.

\> To insure that the breakfast maker problem was not entirely trivial,
\> we implemented a STRIPS formalization of a subset of the domain and
\> tested it using the SNLP non-linear planner [ref] on the omelette
\> making subproblem...
\> ...(The planner ran for 6 hours on a Symbolics XL1200
\> before exhausting its paging disk). The planner was able to make
\>toast
\> and set the table however.
\>
\>If this is what is meant by "too much machinery," I get the point. 6
\>hours!

I think I see why the planner died. At least part of it is that Toast is implicitly told that any pat of butter is the same as any other pat, whereas the SNLP planner was not.

Why are there 15 pats of butter? Why so that the planner must consider more than 23 factorial orders of placing butter and eggs in the frypan.

Agre seems to have some good points, but don't swallow his breakfast without a few grains of salt.

... Jeff

--

De apibus semper dubitandum est. Winni Ille Pu

Date: Mon Aug 31, 1992 6:39 pm PST
Subject: Exploration and Reorganization

[From Rick Marken (920831.1930)]

Martin Taylor (920831 17:20) says:

>"All behavior is the control of perception." But is it?

Being the official sloganeer on the net, I would say:

No. "All PURPOSEFUL behavior is the control of perception".

This slogan captures the fact that there are many aspects of behavior that are neither the actions that are used to control nor the outcomes that result from control of perception. My favorite: the position of my body over time as it plummets into the pool from the high dive. Other aspects of that behavior may be controlled, but I can't do much about it's downward acceleration (unless I'm holding a hang glider).

As to your question

>"What is exploration?"

I don't know. I think that it is different than reorganization because it is systematic; so it is a controlled behavior (some- times). That is, I think it is a program level perception that is carried out for some higher order principle. Thus, I don't think that exploration itself is random; reorganization is. Exploration and reorganization can "look" the same; but I think you could tell the difference using the test.

Best regards Rick

****

Richard S. Marken USMail: 10459 Holman Ave
The Aerospace Corporation Los Angeles, CA 90024
E-mail: marken@aero.org
(310) 336-6214 (day)
(310) 474-0313 (evening)

Date: Mon Aug 31, 1992 7:24 pm PST
Subject: Re: Agre's Cooking problem

[from Avery Andrews 920901.1312]

(Jeff Hunter (920831)

On my reading, C&A's point is not that a certain amount of actual planning is unnecessary, but that what planning there is gets done against a background of routine activity, & lots of experience. E.g. mostly you don't have to figure out the optimal order in which to carry out the breakfast making subtasks, because you've done it lots of times before, and saw people doing it before you had to. So the critique of the planning weenies is not that planning is unnecessary, but that it would be a good idea to find out more about the nature of the routine background against which it happens. In fact, it seems to me that actual planning by humans is incredibly bad when the background is deficient.

Avery.Andrews@anu.edu.au

Date: Mon Aug 31, 1992 11:34 pm PST
Subject: Introducing Myself

Hello All!

My name is Ray Jackson and I'm the Manufacturing Training and Education Manager at the Motorola Computer Group in Tempe, Arizona. My background is in education and, like all good teachers, I've always been fascinated by the psychology of learners. I'm finishing up my Master's at Grand Canyon University where I work closely with Tom Hancock, who introduced me to PCT (as a science beyond Glasser). I've also been in touch with Ed Ford. Tom and Ed introduced me to this network, which I've been on for about two weeks. It is really a wonderful opportunity to get a window into what gets bounced around out here.

For the past two years, my work at MCG is primarily cultural change through education. What I'm finding is EVERYONE in the country is trying to change their "corporate culture", but few do it successfully because they ignore the psychological aspects of empowerment. As behaviorism is the engine for autocracy, PCT provides an understanding of empowerment that goes far beyond the widespread superficial notions. Failing to recognize this does nothing more than alienate people as they go through the motions of empowerment within the autocratic mindset.

All my internal consulting, including the formal teamwork curricula I've contructed, is firmly rooted in PCT - inasmuch as I can understand and articulate it to the leaders and team members in the plant. For the most part, though, I will admit that my PCT knowledge is mostly intuitive. But my formal understanding is growing daily, and I hope soon to be an active contributor on the net.

Finally, Bill, thanks for your vision and your accessiblity. Happy belated birthday, as well. Hope to exchange views with you soon.

Best to All, Ray

****

Ray L. Jackson 602-963-6474

3613 W. Saragosa St.
Chandler, Az 85226
attmail.com!rljackson