Date: Mon Jun 01, 1992 7:24 am PST
Subject: Thanks, penni

[From Bill Powers (920601.0900)]

Penni Sibun (920531) --

The language of the "pengi" article was a bit strange to me -- I'll give it a closer look. I don't recall seeing anything in it about behavior controlling perception, but that concept may be hidden in it. I certainly agree about making choices -- I don't think "decisions" play a very large part in behavior except when we get into internal conflict.

Your precis of the main idea is interesting. We, of course, have been fighting against the idea of behavior that you describe (plan then execute) for a very long time, and control theory is almost entirely about systems that live in a real-time world and interact with it as it changes. Is there any concerted move in AI away from the old view?

Thanks for the updated addresses: I sent my post to the MIT address that appeared on the paper. I'll resend it right now.

What's "imho?"

Best Bill P.

Date: Mon Jun 01, 1992 8:00 am PST
Subject: Changing k.f

[From Bill Powers (920601.0930)]

Rick Marken (920531) --

Clever thought about higher levels increasing loop gain by increasing k.f, the external feedback factor. In manufacutring, k.f is called productivity.

Assuming that the perceptual sensitivity remains the same (no dark glasses) and that the reference level doesn't change, the result of an increase in k.f should be a decrease in effort, shouldn't it? If the feedback effect is increased, the same output effort will produce too much input. So the output drops.

This shows that something is wrong in our economy. American productivity is the highest in the world (although it's not increasing the fastest). So the amount of goods produced per unit labor is higher than it's ever been. But now we have to have two adults working full time per family to bring home enough money to maintain a standard of living that everyone agrees has dropped. It looks as if something is draining the product of our labors away before we can get it. The consumer's k.f (input of goods per unit effort) has become SMALLER in spite of an increase in the producer's k.f (revenue per unit wage paid). There's a glitch in the system somewhere.

Best, Bill P.

Date: Mon Jun 01, 1992 9:15 am PST
Subject: Re: abstract,concrete,HPCT

[From Rick Marken 920601 09:00)]

penni sibun says:

> the paradigm used to be that
> an agent selects a goal, builds a plan out of actions to achieve the
>goal, and then executes the actions as specified by the plan.

Yeah. But that's old hat. Now it's dynamical systems theory -- point attractors and all that jazz. AI people come up with new ways to be wrong faster than we can say "control of perception". Thus, our criticisms of AI and "action theory" tends to be two years out of date all the time.

> the deeper points of the work--such that an agent and
>its world are mutually constructed and that it follows from this
>mutual construction that the world constrains the choices an agent
>has at any point, so that most of the time deciding what to do just
>isn't a big deal (agre calls this ``leaning on the world'')--have been
>largely missed.

I can see why. What does it mean that an agent and its world are "mutually constructed"??? How does one put this into a model that actually behaves in the real world?? I have a feeling that the c&a paper is about dynamical systems again. I bet those constraints imposed by the world on an agent are like point attractors. I see the "mass spring" model of purposeful behavior lurking in language like "leaning on the world". Maybe you could just give a quick summary of how c&a would model some simple, purposeful behavior -- like pointing at a moving target or taking a sip of tea?

Best regards Rick

Date: Mon Jun 01, 1992 9:41 am PST
Subject: Re: Re: abstract,concrete,HPCT

[From Oded Maler 920601]

[From Rick Marken 920601 09:00)]
>I can see why. What does it mean that an agent and its world are "mutually
>constructed"??? How does one put this into a model that actually behaves
>in the real world?? I have a feeling that the c&a paper is about
>dynamical systems again. I bet those constraints imposed by the world on
>an agent are like point attractors. I see the "mass spring" model of
>purposeful behavior lurking in language like "leaning on the world".

Rick, you are so predictable.. I'm sure that at some level in your hierarchy A&C and Beer evoke the same percepts. Anyway your feeling is incorrect and their work is not using the buzzwords of dynamical systems.

>Maybe you could just give a quick summary of how c&a would model some
>simple, purposeful behavior -- like pointing at a moving target or taking
>a sip of tea?

Since in the sequence/program level your favorite theory still does not have muc
to offer beyond qualitative hand-waiving, maybe don't be so quick in dismissing others' attepmts to approach such problems.

Best regards --Oded

Date: Mon Jun 01, 1992 9:42 am PST
Subject: Pengi model

[From Bill Powers (920601.1000)]

to David Chapman, copy to CSGnet --

OK, Penni Sibun has now opened my mind to your "Pengi" article, and it's obvious that our interests are very close together indeed.

One of my beefs with the AI approach and others (like the "motor program" approach) has been the conception of control as a planning-execution process. Like you, I don't doubt that such things happen, but I agree with you that they can't happen in real time. In fact behavior organized that way doesn't really work very well. The real environment is too dynamic and too full of unpredictable disturbances. The best-laid plans of man gang usually agley. Plans are made under one set of conditions, which are only a snapshot of a changing world, and are executed under another.

One of the problems I've seen in conventional behavior modeling is that the modeler fails to take the point of view of the behaving system. The modeler knows the properties of the environment and of the system being designed, and in creating the desired behavior travels freely back and forth across the boundary between Inside and Outside. So if the model doesn't "do" quite the right thing to the external world, the modeler steps inside the behaving system and tweaks it (increase the gain of this circuit a little, add some compensation to that one), watching the results, until the outcome is right. This, of course, lets the modeler do things that the system by itself could never do. The modeler is actually BEING a whole lot of functions that belong in the model. So most modelers are really trying to model something far more complex than they realize, and are giving their models far more external help than they know.

It's OK for the modeler to watch the results and tweak the model. But the tweaking has to be in terms of the capabilities of the model itself -- that is, the model has to be able to accomplish the result without help. "If I were this model, knowing only what it can know and being able to produce only those outputs it can produce, how would the problem look to me and what would I have to be able to do so solve it?"

I think your Pengi model is coming much closer to my (idealized) view of modeling than to the conventional get-the-job-done one. Pengi works on the basis of perceptions of the current environment ("indexical-functional aspects" for crissake). The information it uses is a representation of the current state of affairs (I call it a perception, you call it "registering.")

I think that to some extent you're still using too much of what YOU know about the situation ("it is both vulnerable (if the penguin kicks the block) and dangerous (because it can kick the block at the penguin)". The penguin can't behave "because" the bee is vulnerable or dangerous, unless you've given it the ability to appreciate such abstract conditions as vulnerability or danger. Those two judgements are outside the universe of bee and penguin and iceblocks. But I think that you're working yourself (as of 1987, of course) away from the third-party approach and toward what I think of as real modeling, because for the most part the basis of behavior in Pengi isn't such abstractions, but real-time interactions with the environment according to a few general and simple rules.

One problem in talking about control theory (my version) with people familiar with currently-popular approaches is that most people who use the words "control system" aren't really talking about control systems, but about S-R or planning systems. Real control systems don't respond to stimuli and they don't plan. They don't precalculate outputs that will have desired results -- that concept has grown up mostly without benefit of any experience with real control systems. Just thought I'd warn you, in case you were identifying "control system" with some other current concept (like Rodney Brooks').

I sent the first communication to your old MIT address, which was on the paper. Then I sent it again when I found the Stanford address. And now this. I'll drop it here and wait for your response, if any.

Best, Bill Powers

Date: Mon Jun 01, 1992 10:39 am PST
Subject: Abstract, concrete, HPCT

Hi, thank you for your messages. HPCT sounds interesting.

> How is "Agre" pronounced?

``Eigri,'' in continental orthography.

> Control theory says that control systems VARY their actions in order to CONTROL their inputs. Not their outputs. What others see as controlled output -- as behavior -- is really just an indirect effect of controlling perceptions. Another way to say this is that control systems control OUTCOMES rather than MEANS. This is why some of your buddies at MIT are on the wrong track: they're trying to build models of motor behavior that specify outputs, where the real system works by specifying inputs. They're forgetting that between muscle tensions and their final effects are many other unpredictable influences that also contribute to the outcome.

You might actually find it useful to look at the work of Chris Atkeson, a roboticist at (I'm afraid to say) MIT, and his students, particularly Eric Aboaf. They make this same point, and have a model of skill learning that sounds similar to what you are suggesting here. (Unfortunately I don't have any cites other than MIT tech reports; you could write to Atkeson at cga@ai.mit.edu.)

OK, Penni Sibun has now opened my mind to your "Pengi" article,

A longer and perhaps clearer exposition of this stuff is in an MIT Press book (``Vision, Instruction, and Action,'' 1991).

knows the properties of the environment and of the system being designed, and in creating the desired behavior travels freely back and forth across the boundary between Inside and Outside. So if the model doesn't "do" quite the right thing to the external world, the modeler steps inside the behaving system and tweaks it

The closest work in AI to addressing this issue is the literature on ``temporal difference'' learning. I've done a litle work in that area with an eye to improving Pengi-like models in the direction you suggest, without any spectacular results so far. You might want to look at the current issue of Machine Learning Journal, which I believe is a special issue devoted to TD techniques. TD, btw, is pretty closely related to both dynamic programming and classical control theory.

David

Date: Mon Jun 01, 1992 12:54 pm PST
Subject: Predictabilty

[From Rick Marken (920601 12:50)]

Oded Maler (920601) says:

>Rick, you are so predictable.

Of course. Should I try to be unpredictable? This strikes me as a compliment rather than a complaint -- though you seem to be saying it as though it's a problem. Was there something wrong with Galileo's predictable response to philosophers who argued that the world must be stationary. Is there something wrong with my daughter's predictable response to my claim that 3 times 3 is 10? Yep, I'm happily predictable.

>I'm sure that at some level in your
>hierarchy A&C and Beer evoke the same percepts.

Yes. I perceive thesed models as OPEN LOOP. If I'm wrong, that's fine with me. So far (after 20 years of dealing with these kinds of models) I have not been wrong yet -- such models are either designed to be as open loop as possible or they deny the significance (or existence) of those aspects of themselves that actually are closed loop. So far, I have found only PCT to be based on the observation that organisms control and that what they control is their own perceptual experience.

> Anyway your feeling
>is incorrect and their work is not using the buzzwords of dynamical
>systems.

Just a guess based on minimal data. It sure sounded like they were relying on physical constraints to solve the "planning" problem. If not, fine. And if the model is closed-loop control of perceptual input variables, so much the better.

>Since in the sequence/program level your favorite theory still does not
>have much to offer beyond qualitative hand-waiving, maybe don't be so quick
in
>dismissing others' attepmts to approach such problems.

I didn't mean to seem dismissive, though when it comes to being dismissive I think PCT has more often been the dismissee rather than the dismisser. I've tried to understand and build models based on a lot of behavior theories besides control theory. I have yet to run across one that actually tries to model control. Thus, I have yet to find one that is architecturally equivalent to PCT (ie. it is build around the control of perceptual variables).

I don't know that PCT types have been hand-waving when it comes to control of higher level variables -- we admit that we don't know how to do it yet. (In order to do it we have to learn how to build systems that can perceive the degree to which a sequence or program is taking place. Right now I think the idea is to walk -- literally -- before we can do calculus problems.
I think the "little man" model -- which produces realistic behavior in a realistic environment -- is far more impressive and contributes more to our understanding of the organization of behavior than models which produce complex sequences and programs that only barely resemble what people do -- and in unrealistic environments to boot. I imagine that is a predictable preference). I do assume that control of sequences and programs, etc is based on the same principle as control of other variables -- it is control of perception. If this is true, then the approaches that you say I should not dismiss must be wrong. The only way these approaches could be right is if the sequences and programs are not controlled variables. For the sake of all those people working real hard on open loop models of sequential and programmatic behavior, I hope they aren't.

Predictably best regards Rick

Date: Mon Jun 01, 1992 12:59 pm PST
Subject: Re: Thanks, penni

[From Chris Malcolm]

I recommended Bill's "Behaviour as control of Perception" book to Brooks and his group at MIT in 1987, and some of them later discussed it with me and found it interesting. I can't remember if I ever discussed it with Chapman, but I certainly have discussed it with Agre (who did his PhD with Brooks) and briefly sojourned in Sussex. Nobody else, however, can speak for Agre's opinion of anything :-) I mention it to students in my robotics course here at Edinburgh (some of whom now post to this group), and it features in current robotics research at Aberystwyth University. I first heard about it from Karl Kempf in about 1982, lately of Intel and not concerned with robots now, but once leading McDonnel Douglas's assembly robotics research programme. He said it was popular among roboticists who were interested in learning robot design architectures from nature.

Date: Mon Jun 01, 1992 1:09 pm PST
Subject: U.S. economic "glitch"

From Greg Williams (920601)

Regarding Bill Powers' comments about a glitch in the economy, I think it goes back to 1974 and is called oil prices. Currently, the cost of oil is about four times what it was in 1973, and a rough average for cost of goods (excepting electronics) is about times two -- which would require TWO breadwinners instead of 1973's one per family to keep maintaining about the same standard of living. The REALLY interesting point, for me, is that the pre-73 standard is VERY STRONGLY being maintained. Looks like an importnt reference level: don't let your living standard drop, even if you have to work twice as hard. Quite a high loop gain, I suppose!

A couple of years back, Lester Thurow (sp?) -- MIT economist -- claimed in THE ATLANTIC that he was "puzzled" by the decline in rate of growth in American productivity starting in 1974. If you were buying new machine tools for your business and they suddenly went up in price by a considerable %, I guess you would curtail such productivity-enhancing purchases a b. Why don't economists EVER seem to think in terms of reference levels????

Best, Greg

Date: Mon Jun 01, 1992 1:33 pm PST
Subject: On again

From Tom Bourbon [920601.1]

It took far less time than anticipated for us to get started on the system here in Galveston. Mail to Andy Papanicolaou or me can be sent to: PAPANICOLAOU@UTMBEACH.BITNET.

Date: Mon Jun 01, 1992 3:02 pm PST
Subject: U.S. economic "glitch"

[From Rick Marken (920601 15:15)]

Greg Williams (920601)

Great post Greg!!!

By the way, in my loop gain post yesterday I said the units wrong.

If K = k.o * k.e

Then k.o is in units of output/input; k.f is in units of input/output.

I said it backwards before.

You say:

> The REALLY interesting point,
>for me, is that the pre-73 standard is VERY STRONGLY being maintained.

Yes!

>Looks like an importnt reference level: don't let your living standard
>drop, even if you have to work twice as hard. Quite a high loop gain, I
>suppose!

Yes indeed. I think I've been willing to perceive a slightly lower standard of living than that controlled by my parents -- but not much lower. If it wasn't for the microwave oven (to increase our k.f) I don't think we could have generated enough k.o to maintain our perception at what I consider our relatively humble (by LA standards) standard of living. That new light bulb could really help things.

Best regards Rick

Date: Mon Jun 01, 1992 3:02 pm PST
Subject: correction

[From Rick Marken (920601 15:30)]

I wrote k.e when I meant k.f. So it should read:

K = k.o * k.f

Then k.o is in units of output/input; k.f is in units of input/output.

This makes me realize that the main source of disturbance to high level variables is probably the system itself! Maybe it's because I make and then correct so many mistakes when generating intended sequences, programs, etc that I am acutely aware that these must be controlled variables -- and not generated outputs?

Best regards Rick

Date: Mon Jun 01, 1992 3:04 pm PST
Subject: situatedness

[from penni sibun (920601.1500)] (i'm actually not clear on the rules for this notation....--what time zone?)

[From Bill Powers (920601.0900)]

The language of the "pengi" article was a bit strange to me -- I'll give it a closer look. I don't recall seeing anything in it about behavior controlling perception, but that concept may be hidden in it. I

quite possibly not--i was referring to the tenets of the research approach in general.

Your precis of the main idea is interesting. We, of course, have been fighting against the idea of behavior that you describe (plan then execute) for a very long time, and control theory is almost entirely about systems that live in a real-time world and interact with it as it changes. Is there any concerted move in AI away from the old view?

well, there's a little more respect for situated approaches like c&a's, esp. among roboticists and would-be roboticists. in mainstream ai, there are now things like the oxymoronic ``reactive planning'' and ``hybrid planning,'' where peole are trying to build systems that integrate planning and situated activity. of course, this just pushes the interesting work onto deciding when to do which--and in turn, this is an approach that is still putting a lot of emphasis on deciding.

What's "imho?"

it's net jargon (i've not seen it elsewhere) for ``in my humble opinion.''

==

[From Rick Marken 920601 09:00)]

> the paradigm used to be that
> an agent selects a goal, builds a plan out of actions to achieve the
>goal, and then executes the actions as specified by the plan.

Yeah. But that's old hat.

this may seem old hat to you because i was talking about 1986. regrettably, though, some folks in ai are still wearing this hat, or a version of it w/ a reactive feather clipped on.

wrong faster than we can say "control of perception". Thus, our criticisms of AI and "action theory" tends to be two years out of date all the time.

the mit ai lab isn't what it used to be, but do you really think people there could be two years out of date on ai?

> the deeper points of the work--such that an agent and
>its world are mutually constructed and that it follows from this
>mutual construction that the world constrains the choices an agent
>has at any point, so that most of the time deciding what to do just
>isn't a big deal (agre calls this ``leaning on the world'')--have been
>largely missed.

I can see why. What does it mean that an agent and its world are "mutually constructed"???

well, obviously that's a buzzphrase, and i can't expect it to make a lot of sense to you w/o telling you the story behind it. i think i could do a credible job of it in person, but it would take a long time, and i don't have the wrists for doing it on the net. we're also talking about an entire philosophical approach here. it's like my trying to figure out what's *really* going on in pct: i look at most of the messages on here in the apst month and i've talked to avery, but i'm far from convinced of what y'all are doing. these things take time.

> How does one put this into a model that actually behaves in the real world??

good question. it's certainly something *i* would like to do research on.

> I have a feeling that the c&a paper is about dynamical systems again. I bet those constraints imposed by the

you may be right: if you were to make a convincing argument to that effect, i bet a lot of people w/b interested.

> Maybe you could just give a quick summary of how c&a would model some simple, purposeful behavior -- like pointing at a moving target or taking a sip of tea?

no, actually. i encourage you to read their work and if you're interested translate between their paradigm and yours.

==

> OK, Penni Sibun has now opened my mind to your "Pengi" article,

> A longer and perhaps clearer exposition of this stuff is in an MIT Press book (``Vision, Instruction, and Action,'' 1991).

by chapman. i'd recommend this too, though i think the abstract/emergent paper is a quicker introduction. the best intro to the theory is agre's thesis, which unfortunately is only an mit ai lab tech report; when avery comes back on line he can perhaps give details on obtaining one.

--penni

Date: Mon Jun 01, 1992 4:44 pm PST
Subject: Re: Thanks, penni

[From Chris Malcolm]

I recommended Bill's "Behaviour as control of Perception" book to Brooks and his group at MIT in 1987, and some of them later discussed it with me and found it interesting. I can't remember if I ever discussed it with Chapman, but I certainly have discussed it with Agre (who did his

PhD with Brooks) and briefly sojourned in Sussex. Nobody else, however,

no, chapman did his phd w/ brooks; agre's was w/ mike brady.

--penni

Date: Mon Jun 01, 1992 6:08 pm PST
Subject: UnNettable Variables

[From Rick Marken (920601 18:00)]

I just read a post from penni sibun that made me realize how difficult it is to communicate important nuances of speech over the net.

In an earlier post, penni said:

> the paradigm used to be that
>an agent selects a goal, builds a plan out of actions to achieve the
>goal, and then executes the actions as specified by the plan.

And I replied:

>>Yeah. But that's old hat.

And penni replies:

>this may seem old hat to you because i was talking about 1986.
>regrettably, though, some folks in ai are still wearing this hat, or a
>version of it w/ a reactive feather clipped on.

I was just kidding about the "old hat" stuff. Unfortunately, the net did not pick up my tone of sarcasm. I was sarcastic because I think that theorizing in this area tends to be rather trendy; instead of being aimed at trying to understand the phenomenon under study and what might be required to explain it there is a tendency to go with the currently sexy technology. Your description of the "planned output" approach to behavior seemed quite concise, accurate and timely.

I also said:

>>ai is wrong faster than we can say "control of perception". Thus,
>>our criticisms of AI and "action theory" tends to be two years out
>>of date all the time.

and penni said:

>the mit ai lab isn't what it used to be, but do you really think
>people there could be two years out of date on ai?

Again a misunderstanding. I meant WE (PCTers) are always two years behind because there is a new model to compete against whenever we publish a paper that says "reinforcement theory" can't work or "motor control theory" can't work or "coordinative structure theory" can't work. I think these are actually all the SAME theory in different outfits -- at least they are based on the same conception of behavior (that it is generated output). So I am sure that the mit lab and all the other labs are quite au courant. We (PCTers) are the "old hat" ones. Bill Powers formulated PCT in the 1950s and the basic tenets of the theory haven't changed yet -- because all the research supports the theory to the third decimal place. So we're standing here watching all these other colorful characters roll by -- dismissing PCT as "old hat", by the way-- chasing after the latest rainbow. I think I know why this happens -- why people are always chasing hot new sexy theories. It is because they don't know what control is and they are, thus, not trying to explain it. If you don't want to explain control, you don't need to even consider control theory. So it's not a "theory" problem; it's a "phenomenon" problem. AI type theories see behavior as generated output; PCT is based on the observation that ALL behavior is controlled input. If the AI types are right about the nature of behavior, then their search for sexier models of output generation will be satisfying. If the PCT types are right, then AI's search for sexier theories will be as empty and unsatisfying (in the long run) as the singles bar scene.

I don't expect (or encourage) anyone to favor PCT over any other theory of behavior until they understand (and can demonstrate to themselves at any instant) the PHENOMENON OF CONTROL.

Hasta luego Rick

Date: Mon Jun 01, 1992 7:23 pm PST
Subject: Headers; economics

[From Bill Powers (920601.2000)] --

to Penni Sibun (920601) --

Your own local time in the hand-written header. It's just a way of referring to multiple posts by the same person on same or different days. I suppose you could use a,b,c... but the date, at least, helps when looking up an old post. Managing past input on this net is a major problem for me; mostly I don't.

----

to Greg Williams (920601) --

Good point about the reference levels. But I think there's a deeper glitch in the economy than just oil prices.

The problem is that there's a basic conflict between consumers and producers -- the same one that communism tried and failed to resolve. It hasn't gone away. The split between wage income and capital income in the for-profit sector (government is not-for-profit) is about 40/60 -- 40 percent for labor, 60 percent for owners, stockholders, debtholders, etc. This has been pretty close to the ratio since 1930, with the capital-income share having risen slowly from about 53 percent in 1930 to today's approximately 60 percent. The conflict is that receivers of capital income want their share to increase, while wage-earners want it to decrease.

The composite consumer (not the producer) has the reference level of improving the standard of living. This means working fewer hours to obtain ever-better goods and services, or even just to be working and eating instead of not working and not eating. The idea is that technology or ingenuity -- increased productivity -- should be rewarded by obtaining a better life with less prolonged, unpleasant, boring, dangerous, unremunerative, or mind-numbing labor.

The composite producer (with bean-counters in charge) has the reference level of maximizing the return on investment for the owners of the means of production, or those who have invested in it. This means cutting costs wherever possible and charging the most the market will bear for the lowest quality goods or services that can consistently be sold. Cutting costs means, in large part, reducing the cost of labor. When you reflect that cutting material costs is also cutting costs of labor (on someone else's part), it all comes down to cutting labor costs -- if capital income isn't to decrease.

The kicker is that the wage-earners who produce the products have no way of buying the products except with the money they are paid in wages. So if costs are cut by laying people off, substituting cheaper overseas labor, or reducing domestic wages, the result in all cases is that the buying power of the consumers is reduced -- so the goods and services can't be sold at higher or even the same prices, in the same volume. This is where the conflict comes to a focus.

Unfortunately, this system doesn't have any natural reference levels in the middle of its range of operation -- it just has limits. It always tends toward the state where some large number of people is existing at a subsistence level. The only thing that keeps the composite producer from reducing labor costs any further is the fact that a lot more people would begin dying of starvation or untreated illness or would have their physical living conditions reduced to an intolerable state. The result would be an explosion of crime, or revolution. So a balance is reached where the deleterious effects of further reductions in consumer buying power will increase costs (through taxes for welfare) and reduce sales (through loss of buying power) unacceptably. Government tries to alleviate this situation through redistribution -- spending tax money in ways that increases the slice of the wage-earner or dependent. But the composite producer has no such motive, except when so many people become impoverished that the market begins to fall off.

The government and private philanthropies together manage to acquire enough money from the composite consumer to bring the fraction of capital income down to about 40 percent by redistributing income. Evidently, this is the fraction at which the wage-earning or seeking population has to be maintained even to keep the economy in its current state. If there were no redistribution, there is no way that capital income could remain at 60 percent of the total without creating a violent rebellion by starving people.

People talk in the same breath about our prosperity reaching new highs, if more slowly nowadays, and about the increasing split between high-income people and low-income people. The high-income people are also the chief recipients of capital income. They form the high end of the market. So companies who see sales falling off try to aim for the people who have the money: they produce luxury services, labor-saving items and toys, high-tech or disposable goodies, that will attract the small fraction of the population that has the most money to spend. The result, of course, is that the people at the low end find fewer and fewer items they can afford to buy. The people who CAN maintain their 1970 standard of living work like hell to do so (to get to your point). But just in working like hell to do so, they've sunk below that standard of living. And of course, there are far more people who can't get or handle two jobs, who work less than they used to or at lower wages, and are having a more miserable time than ever.

I think that the owners and managers of this economy need a visit from Ed Ford. Somebody has to ask them, "Is it working?" The problem is that their answer is really "yes" -- so far, it's working for them. A CEO earning $3 million per year plus perks can't really complain. But the SYSTEM CONCEPT isn't working for the people who actually make the system go. It's only working for those who own the system or hold its debts.

There is something drastically missing from the hallowed concept of free enterprise. It's keeping the people whom the economy is supposed to serve in the condition of Skinner's rats. This is something that I think control theorists need to be talking about.

Best, Bill P.

Date: Tue Jun 02, 1992 2:44 am PST
Subject: Re: Thanks, penni

[From Chris Malcolm]

penni wrote:

> no, chapman did his phd w/ brooks; agre's was w/ mike brady.

You're quite right, in formal terms this is true, and Brooks was his internal examiner (or maybe external once Brady moved from MIT), but at least latterly, in terms of intellectual affinity, Agre was a great fan of Brooks, they spent a lot of time together, and up to the time that Brooks visited Brady at Oxford (partly for Agre's examination) Brady was profoundly sceptical of Brooks's approach. After that visit Brady softened to saying things like "the subsumption architecture has some important deep ideas, but is very subtle and easily misunderstood", whereas earlier he would say things like "Brooks has lost his marbles."

Date: Tue Jun 02, 1992 7:48 am PST
Subject: Pengi & HPCT

[From Bill Powers (920602.0930)] (copy to Chapman and CSGnet)

David Chapman (920601a) --

>You might actually find it useful to look at the work of Chris
>Atkeson, a roboticist at (I'm afraid to say) MIT, and his students,
>particularly Eric Aboaf. They make this same point, and have a model
>of skill learning that sounds similar to what you are suggesting here.
>(Unfortunately I don't have any cites other than MIT tech reports; you
>could write to Atkeson at cga@ai.mit.edu.)

Is anyone "situated" to check this out? I'm feeling sort of unwilling to go knocking on more doors. Some days the hill looks steeper than others.

> You might want to look at the current issue of Machine Learning
>Journal, which I believe is a special issue devoted to TD techniques.
>TD ["temporal difference" learning], btw, is pretty closely related to
>both dynamic programming and classical control theory.

Ditto.

Back to Pengi for a moment, David. When I was talking about putting the intelligence of the modeler into the model, I was talking about things like this:

"The-block-I'm-pushing
The-corridor-I'm-running-along
...
The-bee-that-is-heading-along-the-wall-that-I'm-on-the-other-side-of."

It seems to me that being able to recognize such things entails a very complex perceptual system, capable of discriminating, recognizing, and naming objects (block), processes (pushing), agency (I'm pushing), relationships (I pushing block), and so on. I realize that we can't model everything -- we have to use black boxes for what we aren't prepared to model yet. But it seems to me that you're doing ALL the work for Pengi instead of just some of it: the perceptions involved are yours, not Pengi's. You're getting simple behaviors out of a system that's basically extremely complex, chock-full of the modeler's knowledge about real-world phenomena. It seems to me that these are descriptions of OUTCOMES of the model's organization instead of descriptions of MECHANISMS FOR CREATING THESE OUTCOMES. Perhaps I'm judging on too little evidence.

A couple of years ago I wrote a program for Clark McPhail (sociology) at U of IL that simulates the movement of actors through a field of obstacles and other actors. Each actor avoids collisions with other objects and actors, seeks a goal position somewhere in the field, and at the same time may seek a particular spatial relationship with another specific actor or a group of actors. This is a simulation of crowd behavior (pardon me Clark, behavior in gatherings). But this simulation doesn't contain the idea of actors or obstacles or collisions or goal positions. Instead, the things a human observer would classify in these ways appear in behavior that is based on very simple control processes. The actors simply sense proximities and adjust speed and direction of motion as a means of controlling several different simple functions of proximities as perceived by each actor.

1. Each actor is equipped to perceive the sum of all proximities to objects on the left, and the sum of proximities to the right (where proximity is an

inverse-square function corresponding roughly to retinal image area subtended by an object).

2. Each actor can sense the direction to the goal (relative to the actor's direction of travel) and the proximity of the goal (same function as above).

3. Each actor can sense the left and right proximity to one specific other actor or the centroid of a subgroup of actors.

4. For each proximity sensed, there is a fixed reference proximity. Deviations from the reference proximities are reduced by changing turn rate or changing velocity. There are four control systems of this kind involved in each actor, of which three can be chosen to operate at the same time.

The result is what appears to be a quite intelligent and complex behavior. The actors thread their way through a crowd of 50 or 100 (up to 255) other actors and randomly-placed obstacles, backtracking out of traps, finding open corridors, and eventually reaching their respective goal-positions without collisions. If they're seeking proximity to another actor, they will follow that actor around the field, sometimes following the same path and sometimes taking short-cuts or running to catch up after a detour around an obstacle or group of obstacles. You can have chains of actors following the leader. Two groups of actors independently seeking different goal positions can thread their way through each other when their general paths cross, simultaneously avoiding stationary obstacles. You get things like the after-you-Alphonse effect.

The behavior of any given actor could be described in words as in Pengi:

The-velocity-of-the-person-I'm-supposed-to-follow
The-rate-of-approach-to-obstacles-on-the-left
The-density-of-objects-in-front-of-me
The-nearest-open-space-in-front-of-me
The-shortest-path-to-my-goal
The-nearness-to-a-collision

And so on. Actually, none of these would be very relevant to how an actor actually works, and most would be misleading (the actors do not detect open spaces in front of them or "densities", although they seem to. Neither are they concerned with "collisions," although a consequence of their organization is that they adroitly avoid all collisions. They do not seek the shortest path to anything, although sometimes they find it). If one were trying to model the observed behavior of these actors, the Pengi-like approach might eventually lead to a similar kind of result, but would do it through an immensely complex system of rules and logic backed up by complex information about the environment, where the "real" actors actually work in terms of a few simple continuous control processes and know only what they can sense of the environment: proximity to the left and proximity to the right.

This is the basic difference I've always seen between the CT approach and AI. In AI, it has seemed to me, the emphasis is on describing appearances from some human observer's point of view, and then doing logical operations on these descriptions. The modeling approach behind CT is to look for simpler underlying mechanisms that would produce the observed appearances, however they are described and with a minimum of influence from the observer's verbal inferences. Complex behavior EMERGES from the CT model, so it can be compared with real behavior for an evaluation of the model. It has seemed to me that in AI, the starting point is the complex behavior, expressed as generalizations intended to INCLUDE the actual behavior in a given instance. That's a very different approach.

Well, it's a shame your facility at Stanford is limited to Unix systems. Some of our computer demonstrations are real short-cuts to understanding our approach to CT, PCT, and HPCT. Some day, perhaps, we will be able to port our demos to Unix systems, but right now most of us work with IBM clones or Macintosh systems (mostly IBM), because we don't have any institutional backing for the most part (I don't), and who can afford a Unix workstation anyway? Not that I'd really want one -- I've been spoiled by Turbo Pascal and Turbo C. Don't you know any teenagers who would let you run some programs on their teeny-weeny DOS machines?

Best, Bill P.

Date: Tue Jun 02, 1992 10:11 am PST
Subject: economics, behavior

[From Rick Marken (920602 11:00)]

Bill Powers (920601.2000) says:

>The problem is that there's a basic conflict between consumers and
>producers -- the same one that communism tried and failed to resolve. It
>hasn't gone away.

>There is something drastically missing from the hallowed concept of free
>enterprise.

Very interesting post. I'm not sure I agree with all the assumptions but much of it seemed quite sound. I just wonder how PCT could help. Coersive approaches (like communism) are definitely ruled out; certainly the one thing PCT says for sure is that efforts to control people -- ie. make them behave according to someone's reference for "proper" economic interaction -- are bound to fail; and, indeed, there is now considerable evidence that coersive economic systems do fail. So what does PCT have to say that can make things better, other than "don't be so greedy"? Economic analysis based on control theory could be most interesting; perhaps you could suggest some approaches to PCT based macro-economic modeling. Don't your assumptions about the reference levels of aggregate consumers and producers violate some of your cautions about statistical modeling?

By the way. I think suggesting to Americans that there might be something wrong with the "free enterprise" system is similar to suggesting to Catholics that there might be something wrong with the church's injuctions against birth control and abortion. I think that no matter how much pain is being created by a system concept, people who believe in it would rather defend it (and die or allow others to die in the process) than try something that might work. That's why I don't imagine PCT can do much to make the world better; at best, it can just show why we are likely to continually cycle through periods of catastrophe and stability. But my prediction is that, if you get any responses to your "economy" post they will be on the order of "the free enterprise system MUST be right".

Bill Powers (920602.0930) says:

>This is the basic difference I've always seen between the CT approach and
>AI. In AI, it has seemed to me, the emphasis is on describing appearances

And, I might add, taking the "appearances" at face value. Behavior appears to consist of events that are caused by an actor. We see the actor move the castle (in chess), lift the bucket, order the "sell" or "buy", etc. It looks like these events are caused by the actor himself. Bill Powers' monumental observation is that this is never the case -- these events are the joint result of influences created by the actor and the environment simultaneously. Moreover the influences produceed by the actor and environment are rarely the same -- although the events are often repeated. This means that the actor must be varying his influence on the event to compensate for any change in the influences of the environment -- a process called control. So the events that look like caused outputs are actually controlled events.

The rest is history. Hasta luego Rick

Date: Tue Jun 02, 1992 1:03 pm PST
Subject: Jargon

[from Gary Cziko 920601.22.30]

I've noticed a bit more jargon showing up on the net lately, stuff like "imho" standing for "in my humble opinion."

While I have no intention of changing anyone's communication style, I would like to remind CSGnetters that a quite significant proportion of our subscribers are not native speakers of English (indeed, this net is quite international, as it should be) . I am continually impressed by how so many of these non-native English users are able to communicate so well in English and I would hope to keep CSGnet discussions as understandable for them (as well as for us native English users) as possible.

On the other hand, I kind of like the little icons like :-) although it did take me a while to figure out that this was a smiling face sideways. Perhaps someone (Bruce Nevin?) could post a list of these and what they mean (compare :-p). These might substitute for the loss of intonation which caused Marken's sarcasm to be misinterpreted.--Gary

Date: Tue Jun 02, 1992 1:12 pm PST
Subject: Technology as Loop Gain

[from Gary Cziko 920602.1500]

Rick Marken (920531) said:

>So it seems that higher level systems can serve to increase
>the loop gain of lower level systems without necessarily changing
>the physical construction of the control system itself. This
>might be a way of looking at the evolutionary advantage of having
>higher level control systems; its a way to increase loop gain
>without too much structural change in the organism. That is,
>higher level systems let you increase K by increasing k.f rather
>than k.o.

This post got me thinking about technological evolution and how it fits into control theory thinking.

The development of technology might usefully be seen as ways of increasing loop gain in the person-environment loop. Rick mentioned increasing the environment feedback factor k.f. The most obvious examples of this are machines that allow us to achieve very large environmental effects with little output--an individual can dig a trench with a back hoe by just twiddling some levers and pedals for an hour instead of breaking his back with a pick and shovel for a week.

But there is also the perceptual side as well. We increase perceptual gain by using instruments like microscopes, cloud chambers and Hubble telescopes.

Put the two together and you get a tremendous increase in environmental control because of the high total loop gain. And you don't need much k.o at all. Indeed, k.o (muscular strength) probably tends to drop as technology drives up k.f and k.i.--Gary

Date: Tue Jun 02, 1992 5:41 pm PST
Subject: smiling w/ a grain of salt

(sibun 920602.1800)

> I've noticed a bit more jargon showing up on the net lately, stuff like "imho" standing for "in my humble opinion."

every subculture has its jargon. i don't think jargon use in its context should be proscribed (esp. if there is ample opportunity to get it explained). oddly, i don't think i've ever used ``imho'' before, but it was the expression that fitted what i wanted to say. now y'all know it ;-}.

the following smiley file i came across in 1987; i've seen ones since then. smileys are basically a free art form: the only real rule is that if the character used for the mouth ``turns up'' it expresses positive affect and if it ``turns down'' it expresses negative affect. most of the eg's in the file i've *never* seen in use. (using nonalpha characters for emphasis is net jargon too, w/ a lot of variation, and downright confusion bet. emphasis and, eg, titles.) this file also doesn't include many plausible smiley variants--such as my own!

--penni

==

Ever wonder about those strange punctuations embedded in some netnews postings?

(hint- turn your head sideways if you don't get these....)

/*---*
/
The Last Whole Smiley Face catalog :-)

(initially courtesy of Symbolics, Palo Alto)

```
   :-)       yer standard smiley face
   :->       same
   : )       tongue in cheek
   :-(       yer standard frowning face
   :-<       same
   :-(*)     submitter is sick of recent netnews and is about to vomit
   :-o       submitter is shocked / submitter is Mr. Bill
 ~ :-(       submitter is particularly angry

   8-)       submitter wears glasses
   B-)       submitter wears sunglasses
   [:|]      submitter is a robot (or other appropriate AI project)
   :>)       submitter has a big nose
   :<|       submitter attends an Ivy League school
   :%)%      submitter has acne
  =:-)       submitter is a hosehead
   :-)8      submitter is well dressed
  8:-)       submitter is a little girl
   :-)-}8    submitter is a big girl
   %-)       submitter is cross-eyed
   #-)       submitter partied all night
   :-*       submitter just ate a sour pickle
  -:-)       submitter sports a mohawk and admires Mr. T
   :-'|      submitter has a cold
   :-R       submitter has the flu
   :-)'      submitter tends to drool
  ':-)       submitter accidentally shaved off one of his eyebrows this morning
   0-)       submitter wearing scuba mask
   P-)       submitter is getting fresh
   |-)       submitter is falling asleep (or is chinese)
   .-)       submitter has one eye
   :=)       submitter has two noses
   :-D       submitter talks too much
  O:-)       smiley face with halo- submitter is acting very innocent
   :-{)      submitter has moustache
   :-)}      submitter has goatee/beard
   :-d~      submitter smokes heavily
  Q:-)       submitter is a new grad
```

```
   (-:       submitter is Australian
   M:-)      submitter is saluting (symbol of respect)
   :-P       submitter is sticking out tongue (symbol of disrespect)

   8:]       submitter is a gorilla
   8)        submitter is a frog
   B)        submitter is a frog who is wearing sunglasses
   8P        submitter is a bullfrog and it's mating season
   8b        ditto
   |)        submitter is a salamander
   :8)       submitter is a pig
   3:-o      submitter is a cow
   :3-<      submitter is a dog
   pp#       submitter is a cow
   pq`#'     submitter is a bull
   }.        submitter is an elephant
    `\

  +O:-)      submitter is the pope.
  C=:-)      submitter is the Galloping Gourmet
  =):-)      submitter is Uncle Sam
  =|:-)      submitter is Abe Lincoln
   4:-)      submitter is George Washington
   5:-)      submitter is Elvis Presley
   7:-)      submitter is Fred Flintstone
     -       submitter is Helen Keller
    :/7)     submitter is Cyrano de Bergerac
   >:*)      submitter is Bozo the Clown
    #:o+=    submitter is Betty Boop
   _: )      submitter is an Indian
   >>-O->    submitter is Gen. Custer
   8(:-)     submitter is Walt Disney

   >: (      submitter is a headhunter (Amazon style)
 -=#:-)      submitter has wizard status
   (:  (=|   submitter is going to be a ghost for Halloween...
   =:-H      submitter plays for NFL
   (V)=|     submitter is a pacman champion

   M-),:X),:-M  sumbitter sees no evil, hears no evil, speaks no evil

    C):-O
    C):-O
    C):-O
    C):-O   submitter is a barbershop quartet
```

Date: Tue Jun 02, 1992 7:03 pm PST
Subject: Smiling thru; misc

[From Bill Powers (920602.2000)]

Penni Sibun (920602) --

```
(/o^")_    Submitter is taking it lying down. (hint: nose up)
```

I split my sides -- thanks :-(haw)

My objection to jargon -- sometimes -- is that it makes things seem important and new that are neither. Consider "situated" modeling. If I understand this usage, it means models that actually operate in some specific environment instead of in an abstract space. If so, that's the only kind of model I've ever used. I didn't know it had a name.

There must be a word for this: the discovering-you've-been-speaking-prose effect. Maybe the word is "retronym" (analog wristwatch). First it was just modeling. Then people started making abstract models built completely of symbol manipulations. Then they needed a word for models that considered a real environment. So we get "situated models", meaning modeling the way it used to be done? I suppose I've guessed wrong.

Rick Marken (20602) --

It isn't the "free" part of free enterprise that I think has gone wrong. It's the people who use their own freedom as a way of getting control of everyone else. What we need is a theory of economics. I don't mean a better theory -- I mean just a theory (that works). Preferably it will have people in it. We can't fix the system until we understand what's going on.

Best Bill.

Date: Tue Jun 02, 1992 7:32 pm PST
Subject: Re: Smiling thru; misc

(from Penni Sibun 920602.2000)

> My objection to jargon -- sometimes -- is that it makes things seem
> important and new that are neither. Consider "situated" modeling.
> ....
> First it was
> just modeling. Then people started making abstract models built
> completely of symbol manipulations. Then they needed a word for models
> that considered a real environment. So we get "situated models", meaning
> modeling the way it used to be done? I suppose I've guessed wrong.

I think you've guessed right. It's happened cause the ai weenies know *only* about the middle state of affairs. Groundbreaking work in ai in the 70s was *all* abstract symbol manipulation. People got everlasting fame and fortune for writing theses about things like the ``blocks world,'' a ``world'' consisting of a table and some (typically 3) blocks. These entities and their relationships can be expressed in propositional calculus and programs could be given plans for moving around the blocks. (Even doing this sort of thing is provably intractable (which is what chapman's master's thesis was about).) it seemed like a good idea at the time, i guess.

i ta'd and then took an intro ai course at two different schools.

this is the stuff that was presented as ai (in the mid-80s). people in ai needed to be hit over the head. i realized this looks silly from the outside, but each field has its own course of development, and ai has been disadvantaged in that for nearly 20 years it was possible to do brilliantly in the field w/o having any knowledge of anything else--just by writing a snazzy program.

oh, well. i don't tell people i do ai any more (``computational linguistics'' is descriptively adequate) because the glitz is gone and the shallowness has become more evident. oh, and it's not gonna get me a job, either--the world is quite saturated w/ ai types!

--penni

Date: Tue Jun 02, 1992 8:12 pm PST
From: Dag Forssell / MCI ID: 474-2580

TO: Lars (Ems)
EMS: INTERNET / MCI ID: 376-5414
MBX: larsky@it.hos.se
TO: Gary (Ems)
EMS: INTERNET / MCI ID: 376-5414
MBX: G-CZIKO@UIUC.EDU
Subject: Systems Science
Message-Id: 94920603041249/0004742580NA3EM

[From Dag Forssell (920602)]

Lars, thanks for your note of May 29. I'll respond in english this time, so I can copy Gary Cziko.

You are telling me that you plan to write a book about 20 noted system theories for use at Swedish universities in the subject of Systems Science. I can't help but wonder what constitutes a theory. Have you read Thomas Kuhn: "The structure of Scientific Revolutions?" Why are there not 20 theories in chemistry? With 20 theories competing, at least 19 have to be bullshit. (An American expression you should understand as a cattle farmer).

By copy of this post, I am asking Gary to send you the "Boilerplate" suggestion I submitted to Gary several weeks ago in its entirety (unless edited already, of course). (Gary does not have to pay several dollars postage, as I do for a long post. I am not on Internet, but on a telephone company service). It contains a description of our Control Systems Group network and my short literature list plus a demonstration of the phenomenon of control. (You did not answer my $64 question about what a control system controls. It is our experience that most control engineers do not understand that part, and give the wrong answer).

Since you are interested in many theories, you may note the book: "Feedback thought in the social sciences" By Richardson. It covers a lot of ground and represents Powers properly.

With my biases, I would think that you would be better off to write a

Swedish book on Bill Powers' work (Which will prove as revolutionary in the life sciences as Newtons was in the physical sciences), rather than write a book that will be obsolete before it is published, as quickly as the AI crowd discovers that they are failing one approach after another. But that is none of my business, as they say over here.

If you subscribe to this group, you will find a very critical attitude towards a great many contemporary "theories" in Artificial Intelligence, Robotics and Social "Systems" Science. The reasons for this will be clearly spelled out if you ask. You can also download past correspondence. The past weeks have some excellent posts.

Perhaps Gary can download a few other goodies for you to review.

With my best wishes for a successful book!

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

P.S. Would you write me "stora }{|" so I can see how the capital letters are encoded.

Date: Tue Jun 02, 1992 8:22 pm PST
From: Dag Forssell / MCI ID: 474-2580

TO: Hortideas Publishing / MCI ID: 497-2767
Subject: Form Letter
Message-Id: 03920603042230/0004742580NA3EM

Greg Williams, Co-chairman
Hortideas Publishing International
460 Black Lick Road
Gravel Switch, KY 40328

Dear Mr. Williams:

Productive cooperation and individual independence are not incompatible. It just seems that way sometimes in our culture, since it promotes individual competition.

Our understanding of what motivates people is abysmally poor. We lack a good model or "paradigm" to help us understand why people do what they do. As a consequence, we often spend our energies in debilitating conflict instead of productive cooperation.

Without a good model, it is very difficult to comprehend, let alone teach, the Deming Management Philosophy or any other approach to Total Quality Management. It is difficult to design an organization or business system with many persons inter-acting without a good understanding of one person acting.

An understanding of the phenomenon of control provides a fundamental insight that puts the problems that result from human interactions in a bright new light. Because it is basic, it has consequences in many ways in many areas of an organization and the results can be dramatic.

To illustrate: The fundamental insight that the earth is round put the problems of navigation and astronomy in a bright new light in an era when "everyone knew" that the earth was flat. The new insight did not invalidate the common sense observation that the earth appears flat locally, but science moved from a dead end to progress.

Likewise, a fundamental insight into the phenomenon of control does not invalidate any wise common sense observation or practice. It just provides an enhanced understanding of seemingly intractable problems. It provides new diagnostic tools rather than cookbook formulas. Our applications teach skillful use of these diagnostic tools.

The enclosed brochure introduces a leadership program based on a detailed model of how people control themselves; how and why people act. The basic principles can be taught to any attentive person, who can also verify them. People trained in the "hard" sciences appreciate the scientific approach and elegant simplicity. Most people are able to begin applying the principles as soon as they understand the underlying model.

With this understanding you can inform, influence, align and lead people with mutual respect. You can teach better cooperation and effectiveness. Your employees are more effective and more satisfied. The company as a whole responds better to your direction and is more effective. Your company's competitive position is enhanced.

Some people will think that "understanding the phenomenon of control" promises a new way to "control" another person. It is precisely the other way around. We show how people control themselves at all times and how you can work with them when you understand control, rather than inadvertently fight them.

I have 25 years experience in engineering, manufacturing, financial and marketing management. My formal education includes an MBA from the University of Southern California and a Masters degree in Mechanical Engineering from Sweden.

I will be pleased to send you an introductory 39 minute audio tape (with script for reading) which explains the background and scope of our programs. The introduction includes a demonstration which will allow you to determine if your associates recognize control in action.

When you receive the introduction, I think you will find the demonstration both enlightening (startling?) and entertaining. Please feel free to share it with your technical, operations and sales managers at any level for their evaluation. This is a win/win program to greatly increase the understanding and effectiveness of anyone who deals with people in any capacity.

Sincerely,

Dag Forssell

Any comments will be appreciated. Tomorrow, we shall put together 150 or so of this. To be mailed to CEO's in the manufacturing industries. I have a lot of names off a database of California manufacturers.

Bills post to me the other day caused me to think through "expectation management" and strive to spell out as clearly as I can what it is we offer.

>Your area of the CSG archive is prominent!

What do you think I am controlling for?

Seriously, if I should bite the dust prematurely, I would want the group to have my work. Until then, I will of course defend what copyrights are mine so I can make a living. It is quite a lot of fun to be creative with this stuff.

Warm appreciation for you both! Dag
Date: Wed Jun 03, 1992 6:01 am PST
Subject: AI & HPCT

[From Bill Powers (920603.0730)]

Penni Sibun (920602) --

You do my heart good. But I don't think all that AI effort was in vain. It's like building up muscles and skills -- at first, all you do is play, and what the play is about is pretty irrelevant. I think that AI people have developed incredible skills at programming, at handling complex thoughts in an organized way. If we could persuade some of these people to turn these skills to HPCT, nothing will have been wasted.

I feel the same about computational linguistics, Harrisian linguistics, Chomskyian linguistics. All these approaches developed skills. Linguistics, under the surface appearance of studying how people use words in language, is really about how the higher levels of perception and control are organized, and about how independent control systems interact with each other through space and time. I've been stubborn about not considering language as a subject in itself, but as a set of very important clues about how the higher levels of human organization work. Naturally, my linguist friends on the net have resisted this idea, but at the same time they've gone a long way toward putting language into an HPCT framework. When they realize that this is a larger and more important goal, they will shift into a higher gear. Maybe they already have shifted.

It's really unfair to talk about things like this with graduate students or those who are just trying to find a niche within their fields. I always tell such people to get the degree, establish the niche, first, because the other people who have control of your destiny (temporarily) aren't going to understand HPCT heresies. There aren't many places in the country where you can just come right out and be a control theorist, and not be ostracized or starve.

The move of some ai-ers away from abstract symbolism and toward analyzing interactions with real environments is a move toward HPCT. The more people who start worrying about how real outcomes are produced and maintained, the more people who are ready to consider hierarchical control processes. We can't push them into it -- the job is more like digging little channels in a field being irrigated, creating paths where it's easier for the water to flow than in the main channel (10 inches of rain per year, here). Downhill is toward HPCT.

How are you coming along, Penni? Is HPCT starting to look a little more obvious all the time? Watch out. Water doesn't flow back uphill.

Chapman send Agre copies of my posts and I got a nice note from Agre, offering copies of some of his papers. I accepted. It would be very nice to get one or both of these bright guys into our conversations.

Best, Bill P.

Date: Wed Jun 03, 1992 9:26 am PST
Subject: Re: A CALL FOR HELP

```
>
>>From Tom Bourbon [052792 -- 0:46]
>   I need copies of books and computer demonstrations on PCT. . . .
I wouldlike to use material from others -- Wayne Hershberger, Bill
>Powers, Clark McPhail,  Ed.  If you can provide a copy of anything
>you wrote or edited, please send it to me.  (
>   Also, please send the latest versions of any demonstrations or
>modeling projects that you would like me to show.  I will have my
>own over-the-shoulder 386 DOS machine, and the organizers will
>provide a room with a DOS machine and a Mac.
>   Please send all material, by the second week in June, to
>my new address (as of 1 June):
>       Division of Neurosurgery E-17
>       University of Texas Medical Branch
>       Galveston, Texas 77550
```

FROM CLARK MCPHAIL

I have been in Washington, D.C. for ten days and am just now sorting through snail and e-mail. First, I did receive the manuscript draft you sent me and it looks quite useful. I have simply been swamped with end of the semester and then research obligations in Washington. I am now back at the home-stand and working away on book #2, Acting Together: The Organization of Crowds. I will gradually piece together a composite draft of the MSS presentations and circulate it this summer before Durango. Sorry I cannot move quicker but such is my life.

Second, I am awaiting reprints of, "Simulating purposive individual and collective action", the lead article in Social Science Computer Review, 10:1-28, coauthored by McPhail, Powers, and Tucker. This is a full blown account with illustrations of the "gatherings" program. I believe that you have the edition of this program (a.k.a. Crowd) from which those illustrations were printed. If you want another copy I will send one by express mail. I will also send a pre-print copy of the SSCR paper if you wish. Regretably I have little else to send. I think there is too little HPCT in my book, The Myth of the Madding Crowd, to warrant including it in whatever you take to France. Correct me if I am wrong on this point.

When will you return to Galveston? The reason I ask is that I will be in Galveston for a family reunion over the long Fourth of July holiday weekend. Nice timing huh?

Congratulations on your new post, on the prospects for what will prove to be an exciting experience for you and a revolutionary intellectual experience for all those who are fortunate to witness your presentations. You do it superbly well; I only wish that I could be there to witness your performance and your answers to the questions I am sure you will provoke. Best wishes. Bon voyage. Break a leg, (and as ballet dancers always say as they go on stage - merde)!

Date: Wed Jun 03, 1992 11:19 am PST
Subject: AI & HPCT

[From Ray Allis 920603.1100]

Bill Powers (920603.0730)

I have been hesitant to post this, because I'm sure I'm not telling you anything you don't know, but the discussion of AI pushed one of my buttons. I wouldn't want the CSG to trip over the problems AI already has.

>But I don't think all that AI effort was in vain.

I do.

> At the "Dartmouth Conference" in 1956, John McCarthy, Marvin Minski, Nathaniel Rochester and Claude Shannon proposed a study of AI "on the basis of the conjecture that every aspect of learning or other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it"
>
> ("Machines Who Think", Pamela McCorduck, 1976)

Given that goal, the polite judgement would be that the study of AI has successfully proven that original conjecture absolutely incorrect. Maybe there are side benefits in better programming methods and the handling of complexity, but 34 years of "AI" has not yet touched "Intelligence".

The study of AI suffers from a few basic conceptual errors: refusal to recognize deductive "reasoning" is not all there is to intelligent behavior; failure to distinguish between "represent" and "symbolize", failure to grasp the fundamental difference between "analog" and "digital" and confusion of "model" with "simulate" (I call this the "Mathematicians' Mistake").

Artificial intelligence researchers have not and still do not distinguish clearly between these ideas. Careless interchange of the terms can (does) lead to confusion about just what is possible and what is not. Some AI claims can be seen, after thinking about the differences between simulation and modelling, to be just preposterous. e.g the attempt at MCC by Doug Lenat to build an encyclopedic system which will be capable of commonsense reasoning.

There are two very different notions meant by "model" and "simulation". If pressed, most AIers will agree they are trying to _simulate_ intelligence (I contend they really mean intelligent _behavior_). Artificial Intelligence can't tell the difference between simulation and modelling. AI people argue whether it's important that "you don't get wet in a simulated hurricane"! Worse, they usually agree that it doesn't matter!

Model: a physical artifact possessing a subset of the properties of the thing modeled. e.g. model airplane, model rocket. It is an ANALOG of the thing modeled.

Simulation: a mathematical description of some artifact or system. The description can be manipulated by the rules of math or equivalently, deductive logic, with the intention of learning/discovering some properties of the thing simulated. If the system is at all complex, you won't see all the implications without the aid of a computer.

The initial statements for a simulation are premises for a deductive argument. Like all such premises, they should read "IF this statement is true, AND IF this relation holds, THEN the following statement is true. Usually, the critical IF is omitted.

A weather "model" (say of a tornado) involves spinning vats of water, or glass containers of air with smoke trails, or something more imaginative. The modeller tries to include as many of the "relevant" characteristics of the modelled system as possible, and leave out "irrelevant" ones. Telling relevant from irrelevant is definitely an art, and controversial. The model _behaves_ according to laws of physics, *even if we, the modeller, misunderstand or are not aware of those laws*.

Note that a model is an analog, existing in the real world, affected by the real world, in ways not necessarily predicted/able by the modeller. e.g. a wind tunnel model may reveal some effects which are a surprise.

A simulation, being a construction of logical statements, is unaffected by events in the physical world. It is totally deductive; its states are absolutely determined by its form. There are no surprises. A simulation can be made to disclose all its implications, but no more than were built into the starting construction.

A weather "simulation" (say of a tornado) is basically a description of the behavior of a tornado AS THE MODELLER UNDERSTANDS IT. Such a simulation contains statements (assertions) such as "The shape of the funnel is related to the air temperature according to the following function". These statements or equations are the "laws of physics" for a simulation. But they can't be as detailed or comprehensive as real physics.

You, when you program a computer to produce a "little man" on its screen, have constructed a simulation. This is a Good Thing, because now you can work out all the implications of the system of logic. Just remind yourself that you are only making explicit the implications YOU PUT THERE IN THE FIRST PLACE.

A "mathematical model" is really a simulation. A "mental model" is really a concept. Then there is "simulated walnut paneling", which is better termed imitation.

I hope I didn't bore you all too much, but it's been bothering me that "model" and "simulation" are used pretty much interchangeably in some of the discussions, and I see that as one of the reasons AI is such a total failure. I just HAD to say it.

Ray Allis

Date: Wed Jun 03, 1992 1:11 pm PST
Subject: habits

From Andy Papanicolaou [920603 15:31]

It seems that most of the basic psychological phenomena have not been dealt with - not systematically at any rate - in the context of CT. Yet it is obvious that if CT is to be accepted as a general psychological theory, it must deal with all (or most) of these phenomena at least as effectively as the alternative open loop approaches.

From a certain point of view, the task of providing a more satisfactory explanation of psychological phenomena than open loop approaches do, does not appear to be very challenging. yet when one sits down to construct any reasonable and consistent explanation for some of these phenomena, the task proves anything but easy.

Take for example, the phenomenon of skill acquisition or habit formation whereby a sensori-motor event sequence, originally produced haltingly, awkwardly with great effort and requiring conscious attention comes to be performed effortlessly automatically and reliably following a number of trials or repetitions.

An instance of that phenomenon could be the gradual acquisition of the skill to pronounce reliably and correctly a speech syllable consisting of phenonemes, some of which do not exist in the learner's native language, like the monosyllabic French word, tu. This word consists of an initial consonant common to both English and French and a vowel sound which is a phoneme in French, but not in English.

To produce that syllable a specific sequence of partially overlapping articulatory gestures is required constituting a pattern with some invariant features. The English-speaking learner cannot at first produce the sound tu, ostensibly because of the novelty of the required pattern of articulatory gestures. Instead, the learner produces sounds closer to the familiar too or tee, realizing meanwhile that he is off his intended target. After several repetitions and with considerable concentration, he begins to approximate the sound tu until, finally, he can produce it effortlessly and reliably. The fact that he can do so by virtue of having acquired the skill or habit of moving different components (muscles) of his articulatory apparatus in the same specific way suggests (without requiring) that a neuromotor plan has been perfected which, once triggered by the intention to produce the sound tu unfolds effortlessly into the appropriate muscle movements.

Needless to say, the notion of a plan or a set of commands from the brain to the muscles is not a control-theoretic notion at all. This way of producing output is generally inefficient because any unforeseen disturbances will make it impossible for the intended perception to materialize.

On the other hand, some disturbances do indeed prevent the production of the articulatory pattern that results in the intended perception of the sound tu - those that prevent the execution of any of the constituent movements at their appropriate time. The lips, for example, must be rounded or else the entire set of articulatory movements will not result in the intended perception of tu. The range of variations in the shape and temporal unfolding of the pattern is rather limited. The perception tu cannot be achieved unless the same muscles are always used to always do almost the same thing - to produce the same pattern of articulatory gestures.

Again, this fact does not compel us to adopt the notion of the "plan", but it does show why the notion has some appeal, and it does urge one to come up with a CT model for the same phenomenon that would have even greater appeal.

In an attempt to come up with such a model, one encounters several difficulties. Here are a couple: Should one consider as the reference signal the auditory percept tu that the learner heard repeatedly from the teacher? It would seem, at first, to be the sensible thing to do. Yet there are all kinds of empirical findings suggesting that no phoneme can be perceived as such unless some "motor plans" for its production have already been formed even though the learner may not be able as yet to bring his articulation to implement the plant (and produce the phoneme): The native speaker of Japanese cannot hear, cannot experience the difference between r and l when first exposed to these sounds. Slowly, she may come to experience them (once a motor plan for their production is presumably formed) but, instead, she would utter either one of them in the place of the one she is supposed to utter. Finally, she may come to articulate them reliably (once her articulators have been presumably trained to implement the neuromotor commands). We may forego the explanation involving neuromotor plans, but we can hardly forego the indications that before acquiring a skill to pronounce new speech sounds we cannot take it for granted that the learner possesses auditory images of the new sounds to use as reference signals.

Another difficulty: It is a fact that the process of acquisition of the skill of pronouncing unfamiliar speech sounds is greatly enhanced if the learner is given the opportunity to see how the speaker's (the teacher's) mouth forms these sounds. Then the learner imitates the part of the gesture pattern that is visible, and this invariably helps. Now, this fact suggests that, initially at least, it may not be the auditory image of the sound but the visual image of the speaker's articulators that serves as a reference signal. Or that these aspects of both images that are available to the learner serve as reference signals. Moreover, once the imitation of the speaker's (teacher's) visible gestures help the learner to approximate the sound, the afferent signals from the learner's articulators become also available as references.

We can go on and on with difficulties, but at some point, we must begin drawing the outlines of a CT model that simulates the process of this sort of habit formation. If you have any suggestions, Tom and I would surely appreciate them.

Date: Wed Jun 03, 1992 2:35 pm PST
Subject: Re: AI & HPCT

There is something I don't understand. I'm new to this list, and have been noticing a rather strong opposition to "AI" as opposed to "HPCT". I'm not sure I understand the distinction that's being made. Why is HPCT not AI? When you program a little man or some other control system and claim to be studying perception, then as far as I am concerned, you are doing AI, which is a field full of many different approaches, from pure connectionist to pure symbolic to everything in between. I haven't seen anything about PCT that places it apart from AI. What concerns me is that there may be a touch of isolationism here. Sometimes a good idea is slow to gain acceptance because its proponents are too militant in their opposition to competing ideas. Perhaps I'm wrong, but I just don't see the sense in talking about "HPCT" vs "AI," although "HPCT" vs "other competing approaches to
AI" makes sense to me. Isn't this confrontational attitude just inviting rejection from the AI community? I guess I would like someone to explain to me the reason why HPCT cannot be considered as one paradigm within AI.

I don't wish to be too critical here--I'm just making an initial observation. I could be way off base. Is there perhaps a rejection of the principle of machine intelligence within the control theory community? If so, I'm not sure if I see how such a stance follows from control theory principles. So I guess the question I'm trying to ask here is this: "Might an intelligent robot someday be built using control theory principles?" If the answer is yes then HPCT is part of AI.

Allan Randall
NTT Systems Inc.
Toronto, ON

Date: Wed Jun 03, 1992 2:37 pm PST
Subject: Causal Brain & Intentional Mind

[from Gary Cziko 920603.1725]

In a review of two of Jerome Bruner's books in the latest _Educational Researcher_ (1992, May, _21_(4)), David Olson says:

"Bruner makes a strong case for the mind as the proper object of psychology. He is quite right to warn us of the move to identify the mind with the brain and the tendency to blur the distinction by talk of the _mind/brain_. But the relation between the causal brain and the intentional mind remains almost as puzzling to us as it was to Descartes. And the relation between the personal and the social, despite Bruner's efforts, remains puzzling." (p. 31).

Sort of makes you feel privileged to know about PCT, doesn't it?--Gary

Date: Wed Jun 03, 1992 4:01 pm PST
Subject: Re: AI & HPCT

(sibun (920603.1600))

> [From Bill Powers (920603.0730)]
>
> Penni Sibun (920602) --
>
> You do my heart good. But I don't think all that AI effort was in vain.

oh, i agree. i think the problem is more with having gotten stuck in a very bad rut.

> The move of some ai-ers away from abstract symbolism and toward analyzing interactions with real environments is a move toward HPCT.

isn't it really a move toward the same goals, not toward the theory?

> How are you coming along, Penni? Is HPCT starting to look a little more obvious all the time? Watch out. Water doesn't flow back uphill.

well, i still have the view that it all sounds perfectly plausible, but i don't really understand y'all's stand on a bunch of things, so i'm still waiting to see. i must also confess that i'm pretty busy, and haven't been devoting a lot of time to it.

> Chapman send Agre copies of my posts and I got a nice note from Agre, offering copies of some of his papers. I accepted.

yes; i saw that message. i'll just caution you that agre's book has been a year off for four years. more important, it's going to be incredibly dense and inaccessible (at least as of the most recent draft i've seen, and he didn't seem disposed to lightening up much when i talked to him). i still think his thesis (to which the book by now bears little relation) is wonderful, and i hope we'll figure out how to get y'all copies.

> It would be very nice to
> get one or both of these bright guys into our conversations.

well, i'm not sure ``nice'' is the word....the a&c insiders' list is about as different in tone from this one as possible. i was actually struck by this list's ``nice'' tone--it makes me feel uncomfortable! (not in a *bad* way--i'm just not used to such cordiality and politeness on the net.)

btw, here's some more info on one of the things chapman recommended:

Date: Wed, 3 Jun 1992 08:25:25 -0700
From: rich@gte.com (Rich Sutton)
Subject: Special issue on reinforcement learning

Those of you interested in reinforcement learning may want to get a copy of the special issue on this topic of the journal Machine Learning. It just appeared last week. Here's the table of contents:

Vol. 8, No. 3/4 of MACHINE LEARNING (May, 1992)

Copies can be ordered from: Outside North America:

Kluwer Academic Publishers
Order Department
P.O. Box 358
Accord Station
Hingham, MA 02018-0358
tel. 617-871-6600
fax. 617-871-6528

Kluwer Academic Publishers
Order Department
P.O. Box 322
3300 AH Dordrecht
The Netherlands

==

[From Ray Allis 920603.1100]

i agree that ai types are pretty confused about what it means to model something. (perhaps tellingly, i never hear them talking about *simulating* something.) but i diagree w/ your precise distinction, viz.,

> Model: a physical artifact possessing a subset of the properties of the thing modeled. e.g. model airplane, model rocket. It is an ANALOG of the thing modeled.
>
> Simulation: a mathematical description of some artifact or system. The description can be manipulated by the rules of math or equivalently, deductive logic, with the intention of learning/discovering some properties of the thing simulated. If the system is at all complex, you won't see all the implications without the aid of a computer.

and the conclusion you draw about computers:

> You, when you program a computer to produce a "little man" on its screen, have constructed a simulation. This is a Good Thing, because now you can work out all the implications of the system of logic. Just remind yourself that you are only making explicit the implications YOU PUT THERE IN THE FIRST PLACE.

by your definition, a computer program is *not* a simulation, because it runs on a computer (which is governed by the laws of physics or whatever really runs the universe). your program is ultimately constrained by the computer, not by math/logic. obvious examples are memory size, degree of parallel processing, size of the largest (or smallest) floating point number, and so forth. who knows what else (that is, you can't convince me that the situatedness of the program in the computer doesn't affect the program in ways you haven't accounted for).

==

> From: "Allan F. Randall" <randall@dretor.dciem.dnd.ca>
>
> There is something I don't understand. I'm new to this list, and have been noticing a rather strong opposition to "AI" as opposed to "HPCT". I'm not sure I understand the distinction that's being made. Why is HPCT not AI?

i think we're talking about ``communities of practice,'' that is people who know each other, or a least know about each other's work.

i think it's safe to say only a tiny fraction of the people who ``do
ai'' have ever heard of hpct; in that sense, then, hpct is not ai.

--penni

Date: Wed Jun 03, 1992 9:31 pm PST
Subject: Models & simulations; habits

[From Bill Powers (920603.1700)]

Ray Allis (920603.1100) --

RE: AI's hope of computer intelligence:

>Maybe there are side benefits in better programming methods and the
>handling of complexity, but 34 years of "AI" has not yet touched
>"Intelligence".

Didn't Columbus think that if he sailed West, he would reach Cathay? I
think it has been conclusively proven that the initial proposition was
false -- but he found something else (which a lot of people who live near
here wish he hadn't found, but that's a different subject). I haven't been
impressed by the symbol-manipulations of AI, either, but I sure do admire
the techniques. And some day, the AI types are going to wake up to the fact
that what they've been DOING, themselves, is what they should have been
studying. If you want to understand the organization of human behavior, you
can't dismiss anyone's approach to it: the approach itself is evidence.

> ... confusion of "model" with "simulate" (I call this the
>"Mathematicians' Mistake").

>A simulation, being a construction of logical statements, is unaffected
>by events in the physical world.

>You, when you program a computer to produce a "little man" on its
>screen, have constructed a simulation. This is a Good Thing, because
>now you can work out all the implications of the system of logic.

>Just remind yourself that you are only making explicit the implications
>YOU PUT THERE IN THE FIRST PLACE.

I think all these comments relate to a particular kind of modeling or
simulation, the kind that is based on logical statements or empirical
generalizations. Your version of a tornado model seems to be in the same
vein. But there's a different sort of stimulation that doesn't use any
logical statements and isn't a generalization. A supercomputer tornado
model doesn't just say that a funnel will have a certain shape as a
function of temperature. It doesn't actually deal with tornadoes or funnels
at all. It deals with little packets of air that are subject to laws of
physics. Given a certain water content, density, velocity, and temperature
as a starting point for each packet, the computer simply applies the laws
to generate the next state of each packet, and the next, and so on, one
millisecond following the next. All the packets interact with each other

according to their nature and the laws of physics -- as you say, to the extent that we understand their nature and the pertinent laws. The computer program presents a picture of what all the packets are doing as time progresses. It's up to a human observer to give the result a name, such as "tornado" or "funnel."

If you start with one set of initial conditions, you get a warm summer breeze. Start with different conditions and you get a tornado. The point of this kind of simulation is to see what we can't observe directly: how physical factors influence the results. You can't run a real tornado over and over, varying the humidity a little each time to see its influence. But if you have a good model, you can do this via simulation in a computer.

The Little Man works that way. The basic simulation doesn't use logical statements. It is built from mathematical descriptions of how the parts of the body behave or are guessed to behave. How much signal does the tendon receptor generate under a given stress from a tensing muscle? How much feedback signal is there when a muscle lengthens by a specific amount? How much does the muscle itself stretch when subject to tension? How much shortening is there in the contractile part when signals reach it from the spinal neurones? How does the arm respond physically to couples -- torques -- applied at its joints? The answers to these questions are not in the model; the questions are posed in terms of adjustable parameters. These parameters all have direct physical significance.

By running the simulation and adjusting the parameters to get the closest possible match of the model's behavior with that of a real arm, we can estimate the values of the physical parameters of the real arm. This is the basic method of modeling that is behind all the physical sciences and engineering. In this world, "simulation" simply means constructing a quantitative analog of the system, organized according to a model, and running it.

Andy Papanicolaou and Tom Bourbon (920603) --

>Take for example, the phenomenon of skill acquisition or habit >formation whereby a sensori-motor event sequence, originally produced >haltingly, awkwardly with great effort and requiring conscious >attention comes to be performed effortlessly automatically and reliably >following a number of trials or repetitions.

All right, I'll take it. This looks to me like a control system gradually becoming organized. It's a mistake, by the way, to assume that a behavior that's always performed the same way is always performed by the same combinations of muscle tensions. The opposite is most likely to be true. Most behaviors are named in terms of outcomes, not outputs.

I think the basic problem here is the distinction between "conscious" and "automatic." Habits are carried out without thought. If one thinks of control systems as always involving thought, then of course it seems that habits can't be controlled processes. But nothing in control theory says that control has to involve thought or even awareness. In fact there are at least six levels of control in my model BELOW the levels we would associate with thinking. Even spinal "reflexes" are control systems.

Thus the gradual acquisition of the ability to pronounce the French "tu" can be viewed as the gradual acquisition of the ability to make the correct sound appear in perception (or more realistically, what one assumes is the correct sound). This is certainly not hard to explain in control-theory terms: make the sound you're hearing match the reference sound. This does NOT have to involve awareness, although it probably does during the learning phase.

>To produce that syllable a specific sequence of partially overlapping
>articulatory gestures is required constituting a pattern with some
>invariant features.

According to my linguistics friends on this net, that is not generally true. There are only a few phonemes that are closely associated with specific configurations of the articulators. We can hear the same phoneme with the articulators in a variety of arrangements; even a single speaker will use different articulations to produce the same heard sound.

This says that a given articulator configuration or behavior does not necessarily pin down the sound that will be recognized. It is auditory perception that defines the range of DIFFERENT input sounds that will be heard as the SAME word, correctly pronounced. All that the articulators have to do is produce one of the configurations that will result in a perception within the acceptable range. I'm sure that a Frenchman can say "tu" in a friendly way, seductively, sarcastically, or condescendingly -- and that each way of saying it involves a different configuration of articulators. I'm not at all sure, by the way, that it's necessary to round the lips to say "tu" correctly. I can say "oo" with my lips stretched in a smile. I think I can say "tu" with recognizeable correctness in the same way. Of course that's just how I perceive it.

I'm sure that some motor learning takes place as the new pronunciation is mastered. But what does that mean? It certainly doesn't mean "training the muscles," as some people say. All muscles can do is contract; you can't train them to do anything else. They pull their ends together, and that's about it. What you can train is a lower-level control system, a kinesthetic control system. And since muscles always behave the same way, training a low-level control system is largely a matter of training its perceptual function, so it's controlling the appropriate function of kinesthetic sensations.

So: it isn't generally true that uttering a particular sound is a matter of standardizing an articulator behavior, so it isn't generally true that a "neuromotor plan" (thought of as generation of a fixed output pattern) is necessary to cause a recognizeable particular sound to occur. I think this is finally settled by experiments with disturbances that required DIFFERENT articulations in order to produce the SAME pronunciation. We went through this on the net some time ago with our linguists, and I believe the consensus was that people do not need any practice to say words correctly even when there is mechanical interference with pronunciation. The articulators are simply used in a different way -- just the way required so that a human listener will hear the same, or nearly enough the same, sounds (even if a sound spectrograph says they're different). This phenomenon can't be explained by the concept of a "neuromotor plan." As far as I can see, it can be explained only as control of perception.

>The perception tu cannot be achieved unless the same muscles are always
>used to always do almost the same thing - to produce the same pattern >of
articulatory gestures.

It seems very unlikely to me that some behaviors would be organized according to one fundamental principle, and other behaviors -- of the same kind -- according to a different one. Such a proposition says that in the nervous system, there are neuromotor plan generators that are used for all actions in which it happens that the same motor output always produces essentially the same perceptual result, but that in others, a control system with a perceptual function, comparator, and output function is used -- a completely different architecture.

There is certainly no penalty for using a control system even when no disturbances occur, or can occur. Control organizations are generally far simpler than plan generators. Just look at what the motor program people have to go through with their arm models. They have to ask the nervous system to specify the acceleration, velocity, and position of the arm at every instant during a movement, and to calculate continually the inverse kinematics of the arm in order to turn those specifications into torque commands that produce the right movement (after passage through the FORWARD kinematics of the arm). By the time the late Leonard Bernstein figured out how to give the orchestra the downbeat in this way, the players would have finished the symphony. And a control-system model does the same thing, far better, with only a tiny fraction of the computations.

>Yet there are all kinds of empirical findings suggesting that no >phoneme can be perceived as such unless some "motor plans" for its >production have already been formed even though the learner may not be >able as yet to bring his articulation to implement the plant (and >produce the phoneme): The native speaker of Japanese cannot hear, >cannot experience the difference between r and l when first exposed to >these sounds. Slowly, she may come to experience them (once a motor >plan for their production is presumably formed)

What possible evidence can there be of a motor plan that has been formed but is not yet being implemented? I could easily believe that a person gradually learns a perceptual distinction between "r" and "l" -- for one thing, I could ask the person if they sound different. I could also believe that this perceptual learning would go faster if the person was actually trying to produce those sounds at the same time; part of the perception is the perception of how it feels to say the sound. When you say "No, not 'another,' 'A MMMMMother'," you emphasize the closed mouth; "M" is more a feel than a sound, when being compared with "N".

You can't control the sounds of "l" or "r" reliably if you can't perceive the difference between them -- auditorily and kinesthetically/tactily. If you have the wrong feel going with the right sound, you'll feel the wrong articulation and hear its result as right.

> ... we can hardly forego the indications that before acquiring a skill
>to pronounce new speech sounds we cannot take it for granted that the
>learner possesses auditory images of the new sounds to use as reference
>signals.

Those indications just indicate the theoretical bias of the experimenter -- i.e., the explanation that makes the most sense, given the kind of explanation the experimenter prefers. That's just how it is, with control theorists and everybody else.

Before I took any of those indications at face value, I'd have to see the original experimental data for myself, to satisfy myself that this isn't just another of those statistical experimental facts that's true of some of the people in the experiment and false of the others. The problem with standard psychology is that most of its facts are of that statistical variety, but they're always cited as if they were true for every individual. If a fact isn't true of everyone, how can we use it to justify a model that's supposed to be true of everyone?

I could believe this: that you could teach a Japanese how it FEELS to say 'l' and 'r' by manipulation of the mouth and tongue and by demonstrations using mirrors, so that presented with a printed 'l' or 'r' the person could produce an appropriate articulation, by kinesthetic control. And I can believe that while an English speaker might recognize the resulting sound as correct, the Japanese could still not hear the difference, having learned the feel but not the sound. Would that answer the question you're raising?

You propose something similar, but using vision as the controlled variable:

> ... it may not be the auditory image of the sound but the visual image
>of the speaker's articulators that serves as a reference signal.

I think there has to be another step here: seeing the teacher's articulators, the learner has to imagine how it would feel to make one's own articulators look like that (or better yet actually try it, with a mirror). The difficulty with supposing that what is learned is visual feedback is that normally a speaker can't see his or her own articulators. It isn't the LOOK of the tip of the tongue against the hard palate that creates a 't', but the way it feels inside your mouth, where you usually can't see.

I think that "habit formation" carries a lot of old-fashioned freight. It assumes uniformity of output where in most cases there is uniformity only in outcomes. Control theory works with or without awareness, and whether or not disturbances are present. I think "learning" or "reorganization" handles the situation perfectly well, and that the simplest general explanation is that what is learned or reorganized is a control system, not an output.

Best to all, Bill P.
Date: Thu Jun 04, 1992 1:31 am PST
Subject: AI & HPCT

[From Bill Powers (920603.2330)]

Allan Randall (920603) --

I'm sleepy and and I want to go to bed, but I have to answer your query first if I want to sleep.

>Isn't this confrontational attitude just inviting rejection from the AI
>community? I guess I would like someone to explain to me the reason why
>HPCT cannot be considered as one paradigm within AI.

HPCT has been available to be looked at for longer than AI has been around. Newell, for instance, strongly rejected my modest proposal that his theorem prover had the organization of a control system. "Servo theory has nothing to do with it!" Considering our experiences, CTers have no particular reason to care whether AI accepts control theory or not. Particular people, sure, if they can get along in human company.

Actually, AI, the way it's been practiced, is a subset of HPCT. It's concerned with the way people manipulate symbols according to arbitrary logical rules to produce more symbols. That is certainly a real level of human functioning, one level out of 11 that we think we've identified.

If AI concerned itself a little more with how a string of output symbols can result in muscle tensions that create the described situation, it might have more to contribute. It may be headed in that direction. If so, AIers can help elicidate some of the higher-level functions needed for a complete model; they're smart enough to do it as well as anyone could. But AI has never had a complete behavioral model, and it has not recognized two levels in HPCT higher than symbol manipulation, which most people find pretty obvious once they're pointed out. The MAIN thing not recognized in AI is that commands to act can't produce repeatable results in the real world; they haven't discovered feedback control or the reason why it's needed (except maybe at the symbolic level, and there they deny that it has anything to do with control theory -- or used to). Nothing could be more different from HPCT than a model which claims that we first plan our actions and then carry out the plan. Even when people DO try to operate that way, it doesn't work particularly well.

>Is there perhaps a rejection of the principle of machine intelligence
>within the control theory community?

Not at all. I don't really care if machines can be intelligent, whatever "intelligent" means. I suppose they could. I'm only interested in AI as a source of a model for HUMAN intelligence, which so far I don't think it has provided. That means it hasn't told roboticists what machine intelligence would amount to, either. I think it does deal with a certain aspect of human functioning, as I said above. In fact, it seems to be an EXAMPLE of that level of functioning hypertrophied almost to a nonfunctional degree.

>So I guess the question I'm trying to ask here is this: "Might an
>intelligent robot someday be built using control theory principles?" If
>the answer is yes then HPCT is part of AI.

By that definition, physics is part of bridge engineering. If you're suggesting that roboticists might benefit from applying the principles of HPCT, I couldn't agree more. They would probably end up teaching us a lot.

Before such a robot can be built, however, intelligence must be defined by a study of human nature and a model that captures human organization. Once we have such a model, anyone interested can try to implement it in hardware. But implementing it in hardware isn't the goal of us control theorists, except as demonstrations of principles that apply to human beings. AIers tend to define intelligence in terms of the capacity to manipulate symbols -- which, of course, defines the arena in which they compete with each other, often unpleasantly according to Penni Sibun. HPCT could, if we thought in such terms, define intelligence in a much broader way. Such a definition would include the capacity to understand and carry out principles, and to grasp and maintain system concepts -- neither of which types of perception anyone knows how to put into hardware OR software. It would also include the capacity to generate useful categories, control complex relationships, master complex physical skills, etc. Eleven dimensions, at least, of which symbol manipulation is important but not most important.

Off to bed. Best, Bill P.

Date: Thu Jun 04, 1992 6:26 am PST
Subject: Re: PCT & AI

Re: AI & HPCT

From Tom Bourbon [920604 -- 9:20]

[From Ray Allis 920603.1100]
in a response to Bill Powers (920603.0730)

>Note that a model is an analog, existing in the real world,
>affected by the real world, in ways not necessarily predicted/able
>by the modeler. e.g. a wind tunnel model may reveal some effects
>which are a surprise.

>A simulation, being a construction of logical statements, is
>unaffected by events in the physical world. It is totally
>deductive; its states are absolutely determined by its form.
>There are no surprises. A simulation can be made to disclose all
>its implications, but no more than were built into the starting
>construction. ...

>I hope I didn't bore you all too much, but it's been bothering me
>that "model" and "simulation" are used pretty much interchangeably
>in some of the discussions, and I see that as one of the reasons
>AI is such a total failure. I just HAD to say it.

Ray, you certainly were not boring and you need not apologize for your remarks about the meanings of "model" and "simulation"! The issue you address is an important one and PCT modelers must be careful that others understand what we mean by those terms.

I think the way PCT modelers use the terms avoids the mistakes you identify in traditional AI. A PCT model IS an analog -- we assume that the program statements in a PCT model are analogs of the functional organization of a living system AND OF ITS ENVIRONMENT. The phrase in caps is crucial. At the very minimum, there are TWO system equations in a PCT model, one to characterize interactions among functions and variables in the environment, some of which may vary independently of any actions by the control system, and one to characterize the assumed functions and signals internal to the control system. The "running" of the model is a simulation of the moment-by-moment interaction of the PCT model AND its environment. -- the iterative solution of the system equations, with new and often unpredictable values of environmental variables inserted from simulated moment to simulated moment. When a PCT model, or any other model, is simulated in this fashion, there are OFTEN surprises -- some very pleasing to the modeler, others very unsettling. From my own experience, the first time I simulate the PCT model in a new environment is a lot like the times when my children were small and I watched them try a new skill, or perform in public -- I knew, in principle, what SHOULD happen ....

~~~~~~

>From: "Allan F. Randall" <randall@DRETOR.DCIEM.DND.CA>
>Subject: Re: AI & HPCT

>There is something I don't understand. I'm new to this list, and
>have been noticing a rather strong opposition to "AI" as opposed
>to "HPCT". I'm not sure I understand the distinction that's being
>made. Why is HPCT not AI? When you program a little man or some
>other control system and claim to be studying perception, then as
>far as I am concerned, you are doing AI, ...

>... I guess the question I'm trying to ask here is this: "Might an
>intelligent robot someday be built using control theory
>principles?" If the answer is yes then HPCT is part of AI.

Allan, if everyone who is interested in AI were as encouraging and tolerant as you, the appearance of opposition and militancy that you detect on CSG-L would vanish! No matter how politely and innocently a PCT modeler presents the case for PCT, most advocates of AI who bother to respond claim that PCT is "old hat," as Rick Marken recently pointed out, but most of them simply ignore PCT. In that regard, the reactions to PCT by proponents of orthodox AI are similar to those by most other behavioral and cognitive scientists, whatever their theoretical labels.

On the other hand, I believe most advocates of PCT would agree with your conviction that the control-theoretic model IS a model of systems that behave intelligently. In that sense, the PCT model IS a form of artificial intelligence, but by the lights that guide the AI community, PCT is not part of, Artificial Intelligence: The Movement.

~~~~~~

Best wishes to all -- Tom Bourbon

Date: Thu Jun 04, 1992 6:34 am PST

Subject: Cause of Control Movements

[from Gary Cziko 920604.0915]

Rick Marken:

I've have been greatly enjoying reading through your collection of papers in the recently published _Mind Readings_ book.

One study that I find especially intriguing is "The Cause of Control Movements in a Tracking Task." In this experiment you show that you can get subjects to to respond very similarly (r > .99) in the task by having them do it again with the same disturbance pattern, but that variations in the position of the cursor between the two runs are not similar (r usually < .20). You conclude that "This result seems to rule out stimulus variations as the cause of responses which control (stabilize) the cursor."

This is indeed a very ingenious demonstration. But even after about three years now of studying PCT, what you've demonstrated still looks a bit like magic to me. My brain keep saying to me (or at least part of it): "Surely there must be SOME aspect of the cursor which determines the response. If it isn't the simultaneous position of the cursor, then perhaps it is the cursor's position some milliseconds before or the speed or acceleration of the cursor or SOMETHING which determines response." Maybe you felt the same way which led you to use the word "seems" above.

And if what you say is true, how on earth do all the human factors types continue to see the human operator as a transfer function between stimulus and response? What type of function do they find? Indeed how can they find any function at all if in fact there just ain't none, as your research suggests?

I'm looking forward to your response as well as that of any others who wish to comment on this.--Gary

Date: Thu Jun 04, 1992 7:07 am PST
From: Hortideas Publishing / MCI ID: 497-2767

TO: * Dag Forssell / MCI ID: 474-2580
Subject: your letter to CEOs

Dag, I think your newest version of the letter sounds at least as good TO ME as the previous versions. I'm not competent to suggest how a CEO might "handle" it... You'll just have to send it and see!

Keep up your Good Work,

Greg

Date: Thu Jun 04, 1992 9:14 am PST
Subject: AI, PCT, and HPCT

[From Bill Powers (920604.0800)]

Penni Sibun (920603.1600) --

It's awful, isn't it? Here you are trying to get some work done and CSGnet won't leave you alone.

It occurs to me in the cold (actually comfortably cool) light of day that "AI" (as treated on CSGnet) is just a symbol for something else. Artificial, or machine, intelligence is to me just a way of trying to model intelligence; the fact that a machine may end up doing intelligent things is a side-issue. I can see where a simulation might interest some people BECAUSE it's happening in a machine, but I don't see that as much different from its happening in neurons and muscles. To me, the interesting question is "What's happening?" not "What's it happening in?" "AI" is to me a symbol for people who have naively accepted the mainstream scientific conceptions of what behavior is and how it works, as if those conceptions were facts of nature and needed only to be explained. What's wrong with AI is that it's trying to model things that don't happen. That's why it hasn't got anywhere. Perhaps that's why the people in it get so snappish with each other.

That table of contents Rich Sutton so obligingly copied out and that you so helpfully sent to us is very revealing. Here are all these high-powered people running computer models or doing sophisticated abstract mathematics, and what is the subject? REINFORCEMENT! This seems to happen whenever anyone outside the field of psychology tries to step in and show those dumb bunnies in the soft sciences how to do the job right. The first thing they do is accept at face value the explanations of human behavior that those dumb bunnies thought up.

If AIers want to use HPCT, they have to realize that PCT (without the H) has redefined the problems that they're trying to solve. It says to AI (and most other disciplines) that behavior simply does not work the way behaviorial scientists and biologists have imagined it to work. If you're going to simulate something having to do with human organization, you should make sure first that it really exists. Reinforcement is among the things that PCT can show to be nonexistent: there is no special effect of certain stimuli or objects or events that causes organisms to learn, or to do anything at all. Reinforcement is a total misconception of the relationship between organisms and environments.

Consider another but related subject. What would a person have to believe in order to believe in the plan-then-execute type of simulation? The most obvious belief is that if the brain can calculate just the right output signals, those signals will make the muscles produce the commanded "behavior." It's assumed that if regular outcomes of behavior occur, the motor actions creating the behaviors also must have been regular. This is the same assumption on which all the behavioral sciences were founded, and in which they still believe. But it's false.

This assumption arose long ago from taking something for granted and failing to check it out. In fact, if you send the same driving signals to the muscles twice in a row, you'll be lucky to see any resemblance at all between the behavioral outcomes on the two occasions. The only way to get a repeatable outcome is to work with "preparations." Between muscle tensions and the regular results that are recognizeable as behaviors in a real environment, there are innumerable causal stages . At every stage, beginning inside the muscles themselves, there are independent disturbances and uncertainties and bifurcations that add their effects to the stream of causation. Under the plan-then-execute paradigm, this has to mean that variability must increase as you follow this causal chain outward into the environment. In fact -- and anyone could have realized this at any time if they had just looked instead of assuming -- variability is LEAST at the END of this causal chain, and GREATEST at its BEGINNING. This simple observational fact wipes out any idea of output actions being caused by stimulation OR by serial chains of computation in the brain. So it wipes out the foundations for most of the conventional sciences of behavior. And for most of AI.

If you ask how it can be that variable means produce consistent and often closely controlled and disturbance-resistant outcomes, you end up with control theory. That's the ONLY explanation anyone knows of that works. PCT isn't optional; it isn't just an alternative view. It's the inevitable result of admitting that outcomes are in fact under control (meeting a formal definition of control), and seeing, eventually, that the only way to explain this fact is that the organism is controlling its own perceptions of those outcomes. You can verify that this is how it works six ways from Sunday: once you see the phenomenon that actually needs explaining, there's no difficulty at all in showing that it's real. In fact, most people who finally catch on to the basic principle of control find their own examples and proofs -- they're impossible to miss once you realize what you're looking at.

So when I say that a move toward analyzing interactions with real environments is a step toward HPCT, that is exactly what I mean. If this analysis is done thoroughly enough, the analyst will come up against the fact that regular "behavior" -- defined as regular outcomes of motor action -- is NOT a regular function of motor action. The fact will come out that even in environments full of independent and unpredictable disturbances, in which information available to the senses is totally inadequate for predicting the effects of a given output act, organisms can control outcomes reliably and often with great precision. And unless the analyst commits an act of extreme genius and discovers some totally new way of accomplishing this result, this analyst is going to rediscover control theory. Nobody has ever offered a different explanation that can actually account for these facts.

HPCT, with the H, is an embellishment on PCT. In HPCT there is a place for the kinds of studies that have been going on in AI. But with the knowledge of PCT in the background, those studies would take on an entirely new look: the aim would change as the phenomena to be explained are seen in a new way. Nobody would be wasting time talking about conceptions from traditional disciplines that are based on a completely wrong model of behavior itself -- a model that doesn't even deal with the most funbamental facts of behavior.

Best, Bill P.

Date: Thu Jun 04, 1992 9:16 am PST

Subject: Cause of Control Movements

[From Rick Marken (920604 10:00)]

Gary Cziko (920604.0915) says:

>I've have been greatly enjoying reading through your collection of papers
>in the recently published _Mind Readings_ book.

I love you.

>One study that I find especially intriguing is "The Cause of Control
>Movements in a Tracking Task." In this experiment you show that you can
>get subjects to to respond very similarly (r > .99) in the task by having
>them do it again with the same disturbance pattern, but that variations in
>the position of the cursor between the two runs are not similar (r usually
>< .20). You conclude that "This result seems to rule out stimulus
>variations as the cause of responses which control (stabilize) the cursor."

I wanted to call the paper "The cause of control in a tracking task" but
they didn't understand what that meant; so the title is really wrong.
The paper is not about how movements (of the arm or handle) are controlled;
it's about how control (stabilization against disturbances) occurs. This
was beyond the control theory experts who reviewed for the journal. But,
at least they let it get published -- with that stupid change in the title.

>This is indeed a very ingenious demonstration. But even after about three
>years now of studying PCT, what you've demonstrated still looks a bit like
>magic to me.

It is -- just like the path of light being bent by a mass. But once
you have the right model it all makes sense.

> My brain keep saying to me (or at least part of it): "Surely
>there must be SOME aspect of the cursor which determines the response.

It's the part of your brain that can't think in circles (though you can
think in circles around me). There IS, indeed, some aspect of the
cursor that determines the response -- it is the position of the
cursor. But AT THE SAME TIME the position of the cursor is being determined
by the response. The significance of this fact is hard to "think through"
using our usual "lineal" approach to thinking -- a causes b causes c causes...
You just have to trust the math on this one.

> If
>it isn't the simultaneous position of the cursor, then perhaps it is the
>cursor's position some milliseconds before or the speed or acceleration of
>the cursor or SOMETHING which determines response."

That's why I did the experiment -- ALL these possibilities are ruled out
by the lack of correlation between cursor traces. If, for example, the
ACTUAL cause of the response is the value of the cursor 100 msec before
the current display then this would be true during both runs since the
response is EXACTLY the same both times. So the cursor traces would be the
same (high correlation) on both runs -- but they are not. The fact that

the cursor traces are different means that NO aspect of the cursor trace can account for the identity of the responses on both runs. This just falls out of the feedback equations; remember, output depends on disturbance, NOT sensory (controlled) input. This demonstartion just shows that what the equations say is true; the equations imply a result that is just as magical as the results in this paper. Bill Powers discovered this fact about control and showed that it really happens; I got to show that it really happens in another way -- idential responses produced by different inputs.

>And if what you say is true, how on earth do all the human factors types
>continue to see the human operator as a transfer function between stimulus
>and response?

Why, for that matter, has psychology and the life sciences in general continued to see organisms as transfer functions between inputs (or "commands" or "plans") and outputs while completely ignoring the fact that control systems DON'T WORK THIS WAY? I think you know the answer; did YOU really want to understand PCT at first? The life sciences are made out of people with careers, reputations, etc to protect and families to support. It seems quite understandable to me that they would be reluctant to find out that there is a fundemental flaw in the idea that supports their career, reputations, etc; a flaw that says "everything you have been saying about behavior is completely wrong". Why would ANYONE want to go to the troble to find THAT out? Just nut cases like you and me.

> What type of function do they find? Indeed how can they
>find any function at all if in fact there just ain't none, as your research
>suggests?

Their models work because stimulus response models work, with the proper time damping and gain, in a closed loop. We have already gone over this in considerable grisly detail. Basically, they build models that say

o = f(i) output is function of input

and ignore the fact that

i = g(o) as well,

though they have to hook things up properly so that this second equation holds true.

This "conventional approach" to control ignores the fact that the reference level of sensory input is determined by the organism (a VERY important omission; in fact, just about the whole ball game).

But you might wonder why the conventional control theorists never noticed the fact demonstrated in my experiment (no apparent relationship between controlled input and response). Two reasons: 1) they never looked and 2) if they had looked, they would have probably found far more of a relationship than I found because they usually do their tracking experiments with the subject operating near the limits of control; they use high frequency, high amplitude disturbances so control is POOR. When this is true, more of the variance in the cursor (controlled variable) is caused by the disturbance. So, to the extent that there is any control (and since responses correlate with disturbances, not with the controlled variable) there will be a correlation between response and input. The results reported in my paper depend on the subject being IN CONTROL of the cursor; most of the disturbance-caused variance in the cursor is removed by the repsonses of the subject.

I can assure you, Gary, that there is no trick involved in the experiment you read about; you can go out and demo this to yourself anytime. The magic, I'm afraid, is real. The result can only be understood in terms of the simultaneous equations that describe a negative feedback relationship between an organism and its environment -- the situation that exists for all organisms, all the time. You can't understand it in terms of lineal cause and effect because control does not work that way. The results are only magic from the lineal cause effect point of view -- they are a yawner from the PCT point of view. But if one has an investment in the linear cause effect point of view then these results will evoke responses of disbelief or disinterest; and these have, indeed, been the most common responses to this study.

I think that you DO understand the results, Gary. They are only surprising when you look at them from your old point of view -- and the one that is easier for everyone to use, myself most emphatically included -- the point of view of the lineal cause effect model of behavior.

Best regards Rick

Date: Thu Jun 04, 1992 9:57 am PST
Subject: car event

To: CSGnet people
From: David Goldstein
Subject: a strange incident today
Date: 06/04/92

I was driving to work this morning. I was thinking about the discussion of modeling versus simulation, AI versus HPCT. I started to imagine putting a TV camera on the car which was part of a model designed to drive the car to work. I was wondering how the model would respond to a car which suddenly came close from the side direction. Then suddenly, just at that moment, reality intruded itself. The car on my left suddenly swerved in front of me. I quickly braked and avoided the collision. To the best of my recollection, this is the first time I was thinking about this sort of problem. It was kind of spooky. After the incident, the cars around me seemed to be doing maneuvers which seemed somewhat dangerous. This sense of danger continued for a few minutes until I exited the highway I was on. It was OK after that. Was this just a coincidence? Psychic phenomena? Was I in imagination mode and this played a part in the incident? What are your speculations?

Date: Thu Jun 04, 1992 9:57 am PST
Subject: ai

eric harnden(900604)
please excuse my intrusion at this point, for what might not be a particularly substantive contribution to this discussion, but i feel compelled by my own combination of interests to point out that people who engage in AI are not fools. it is well known that rule-based systems can model (or simulate - the usage doesn't affect the issue) only a limited subset of what might be termed intelligent behaviors. ('intelligence' is a shorthand within the community. no-one pretends that the boxes actually think, yet.) they do, however, provide useful tools for information processing, and in fact are alive and well as pre-processors, frontends, output analyzers, and other adjuncts to advanced work in neural networks. and while some of the more farfetched goals of the field may seem meaningless to those primarily interested in 'deep' or potentially realistic models of human cognition, they continue to provide platforms for highly utilitarian research.

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

Date: Thu Jun 04, 1992 11:03 am PST
Subject: Re: Habits

From Andy Papanicolaou and Tom Bourbon [920604 13:50]

Response to Bill Powers [920603.1700], Re: Habits

The hope behind the selection of this specific topic (acquisition of, or learning of the skill or the habit of producing correctly and reliably new speech sounds) was to elicit specific comments that would help us construct a CT model of the process.

It was not meant to elicit a reiteration of the reasons why a CT model would be preferable to an open-loop "plan" model.

The first specific problems or difficulties that we (Tom and I) encountered in thinking about how to proceed with constructing a CT model of this phenomenon was not the distinction between "conscious" and "automatic". We already know that "habits are carried out without thought" and we certainly do not think that "control theory says that control has to involve thought or even awareness".

The specific difficulties we encountered and we spoke of had to do with what we should consider as the reference signal.

We also laid out the facts of the case as we understand them. Facts that the CT model we hope to construct has to take into account. We still believe that it is a fact that to produce a particular phoneme, a time-varying pattern of articulatory gestures is necessary. This pattern has some invariant features. What we mean by that is simply that to produce a sound that is heard as the French /u/, the lips must be rounded and to produce a /b/ sound, the vocal cords must start vibrating before the lips open to let out a burst of air. If the cords start vibrating after the air release no /b/ can be produced, or if the lips are not closed completely and the air is not suddenly released again, no /b/ will be heard.

We also believe that it is a fact - and we so stated - that although there are invariant features to each articulatory gesture pattern, there is a range of variation as well, which accounts for the fact that the same syllable can be pronounced with different accents, angrily, sadly, etc.

So, we believe that it is generally true (and, hopefully, the linguists will concur or persuade us otherwise) that every time a /ba/ or a /ga/ or a /tu/ is heard the corresponding patterns of articulatory gestures contain a set of invariant features.

The linguists may be also able to tell us whether we should take it for granted that the Japanese possess auditory images of r and l, that is they can discriminate these phonemes even if they cannot produce them. And we are not talking here about a statistically averaged Japanese but about concrete individual Japanese who have not been exposed to languages in which r and l are phonemes. Until the linguists persuade us otherwise we will be cautious and we will not take it for granted that once a Japanese hears /relax/ for the first time he can experience r and l as different phonemes.

We trust it is clear that the questions regarding the Japanese' ability of experiencing r's and l's is directly related to the issue of what constitutes, initially, the reference signal and not to any peculiarities of particular linguistic groups.

We also trust it is clear that we are here concerned with the general issue of skill acquisition and not the instance of acquisition of speech sound producing skills not because we think that these skills are NECESSARILY of a different kind than the skill of typing for example but because they force us to consider issues (like the problem of reference signal) which, in the case of the typing skill may have gone unnoticed.

Finally, we hope it is understood that our search for appropriate reference signals does not include "higher order" references like the "intension to communicate" etc. We wish to focus on how to account for the acquisition of the skill in a situation where the learner simply consents to try and produce a particular sound to the satisfaction of a teacher.

So, to summarize, (1) in an attempt to construct a CT model that would account for the acquiring or learning the skill or forming the habit of pronouncing new speech sounds or "reorganizing" to that effect, one of the initial difficulties we experience is what to think of as reference signals.

(2) Bill's comments helped us realize that there may be no consensus as to what constitutes the set of facts that the model should account for. Hopefully, such consensus is possible otherwise each of us will construct models of a variety of private linguistic facts and worlds.

(3) We wish to reiterate our intention to construct a CT model not a hybrid CT plus "neuromotor plan" model of this process. We hope that the facts, including the existence of articulatory invariances, can be accommodated within such a model because to us also "it seems very unlikely that some behaviors would be organized according to one fundamental principle and other behaviors - of the same kind - according to a different one".

Date: Thu Jun 04, 1992 11:57 am PST
Subject: just the same

[From: Bruce Nevin (Thu 92044 14:38:31)]

Avery dropped by yesterday, and we had lunch and some time to chat face to face. A very pleasant visit. He's off to Philadelphia now.

One of his remarks just before leaving was the supposition that the core of the metalanguage is "in the hardware" so to speak rather than in language--neural mechanisms doing things that emerge as language. Of course this must be so. Equally "of course," important aspects of the metalanguage are in language ("word" is a metalanguage word). The question then becomes, what is on each side of the division.

To avoid some predictable objections, I'll put this in considerably more awkward terms. Read whichever version you prefer.

Words are perceptions, i.e. neural signals entering and leaving elementary control systems (ECSs) in a hierarchical control system. In a way that is yet to be understood in the model, word perceptions are brought into correspondence with non-word perceptions for "object" language, and into correspondence with word perceptions (and other "language-internal" perceptions, for lack of a better term) for metalanguage. So "metalanguage" in this sense consists of neural signals, a subset of the neural signals for language. "Metalanguage" in Avery's sense consists of ECSs (and their I/O functions) controlling the neural signals for language. The question then becomes, what is on each side of the division.

I asked a question some time back that seems to me crucial for this inquiry. That question was: can an ECS control for two perceptual signals being the same? (Say, the recognition of the particular dog that carried off my newspaper yesterday.) Recall that this is the condition for many reductions. A word can be reduced to pronoun or to zero, for example, only under this assertion of "you already know what I'm talking about" sameness.

In Generativist theory, this metalanguage assertion of sameness is carried by subscript letters after the words in question, cute and convenient but implausible. In Harrisian operator grammar, it is carried by a metalanguage assertion. It is useful that many reductions are constrained to act only on word pairs whose sameness can be asserted in this way (whereas indexing with subscripts is relatively unconstrained), but the proposed sources with explicit metalanguage are implausible.

How could ECSs do this? Obviously, they could control metalanguage words that never (or rarely) get spoken, are almost always zeroed. We surely have a relation perception with which we associate the word "same".

But this perception of sameness would apply not to the two occurrances of a given word; that would be a different sort of perception, one of repetition of the word. Nor would it apply to the category perception with which the word is associated (if indeed that is what is going on); that would be a perception of something being of like kind. Rather, it would apply to lower-level nonverbal perceptions satisfying the input requirement of the category ECS. These perceptions might differ in detail (Bowser asleep vs. Bowser wagging his tail and panting to go out, two sides of Mt. Shasta, etc.) The perception of sameness then overrides these differences.

Perhaps some lower-level perceptions are in common, as a basis for recognition (same as remembered individual).

We compare present perceptions with remembered and imagined perceptions. If possible we test our relationship with the individual (barks like Bowser when I call his name--must be Bowser).

Is there any other way that a perception of sameness could arise in the model?

It appears that this entails that perception of an individual perduring through time and across occasions depends on the category level precisely to the extent that language is claimed to depend on the category level. Is this an acceptable consequence?

Bruce
bn@bbn.com

Date: Thu Jun 04, 1992 12:10 pm PST
Subject: Re: AI, PCT, and HPCT

[From Oded Maler 920604]

The "Reinforcement" in the machine learning context is used in a very technical sense without any metaphysical assumption about what "behavior" is and all the rest of the S-R psychology you dislike. You just observe how a function behaves and try to adjust it until it optimizes some criterion. You could apply these techniques in the "reorganization" stage.

Date: Thu Jun 04, 1992 1:58 pm PST
Subject: Re: AI, PCT, and HPCT

[Ray Allis 920604.1200]

> (sibun (920603.1600))
>
> i agree that ai types are pretty confused about what it means to model
> something. (perhaps tellingly, i never hear them talking about
> *simulating* something.)

Most of the people I know understand that they are SIMULATING intelligence or intelligent behavior (or trying to). After all, the rules of the game say this is to be done with digital computers, which can simulate anything. The problem is that everyone seems to think that means "PRODUCE intelligence". They act as if a simulation is equivalent to the thing simulated.

> by your definition, a computer program is *not* a simulation, because
> it runs on a computer (which is governed by the laws of physics or
> whatever really runs the universe). your program is ultimately
> constrained by the computer, not by math/logic. obvious examples are
> memory size, degree of parallel processing, size of the largest (or
> smallest) floating point number, and so forth. who knows what else
> (that is, you can't convince me that the situatedness of the program
> in the computer doesn't affect the program in ways you haven't
> accounted for).

Won't try. I'll agree the 'situatedness' affects the _implementation_ and _operation_ of a program. Which is to say the operation of the physical computer.

But I'd like to point out that I'm not talking about the physical computer, I'm talking about the structure of symbols which the (physical) computer manipulates. (As as side observation, symbol manipulation is what digital computers do; what they are designed to do, and all they can do. The more interesting analog machines are unfortunately unpopular these days.)

A digital computer is indeed a physical artifact, affected by lightning, dynamite, baseball bats etc. But, a program for such a computer is an abstract, non-physical, logical structure of symbols and their inter-relationships. That includes all programs; the operating system, the compilers and the simulation. Lightning et. al. can only affect a _physical implementation_ of a symbolic structure. Think of The Lord's Prayer or a sonnet by Shakespeare. How does the physical universe affect it? A symbolic structure which is a simulation of a human heart is just as unaffected by reality.

> ==

> [From Bill Powers (920603.1700)]
>

> I think all these comments relate to a particular kind of modeling or
> simulation,

I intended to point out that modeling and simulation are fundamentally
different things.

the kind that is based on logical statements or empirical
> generalizations. Your version of a tornado model seems to be in the same
> vein. But there's a different sort of stimulation that doesn't use any
> logical statements and isn't a generalization. A supercomputer tornado
> model doesn't just say that a funnel will have a certain shape as a
> function of temperature. It doesn't actually deal with tornadoes or funnels
> at all. It deals with little packets of air that are subject to laws of
> physics.

No it doesn't. It's not a "supercomputer tornado model", it's a
simulation, and it deals with _symbols_ for little packets of air that
are subject to laws of physics.

Given a certain water content, density, velocity, and temperature
> as a starting point for each packet, the computer simply applies the laws
> to generate the next state of each packet, and the next, and so on, one
> millisecond following the next. All the packets interact with each other
> according to their nature and the laws of physics -- as you say, to the
> extent that we understand their nature and the pertinent laws. The computer
> program presents a picture of what all the packets are doing as time
> progresses. It's up to a human observer to give the result a name, such as
> "tornado" or "funnel."

Well, sort of. There are no 'packets of air', and therefore no
interaction among packets. The computer operates on an arrangement of
symbols, producing another arrangement of symbols. The connection
between either arrangement and 'reality' is entirely up to the
observer; any association between symbol and symbolized exists entirely
in the observer's 'mind'.

>
> If you start with one set of initial conditions, you get a warm summer
> breeze. Start with different conditions and you get a tornado. The point of
> this kind of simulation is to see what we can't observe directly: how
> physical factors influence the results.

But you don't see how physical factors influence the results. You only
see how the changes you made in the logical structure of the simulation
affect the results. The relationship between the physical factors and
the logic is in your mind.

You can't run a real tornado over
> and over, varying the humidity a little each time to see its influence. But
> if you have a good model, you can do this via simulation in a computer.

See above.

> The Little Man works that way. The basic simulation doesn't use logical
> statements. It is built from mathematical descriptions of how the parts of

> the body behave or are guessed to behave.

These are logical statements. Deductive logic.

How much signal does the tendon
> receptor generate under a given stress from a tensing muscle? How much
> feedback signal is there when a muscle lengthens by a specific amount? How
> much does the muscle itself stretch when subject to tension? How much
> shortening is there in the contractile part when signals reach it from the
> spinal neurones? How does the arm respond physically to couples -- torques
> -- applied at its joints? The answers to these questions are not in the
> model; the questions are posed in terms of adjustable parameters. These
> parameters all have direct physical significance.

So you say. That's not meant to be flip, but they only have direct physical significance if you say so. I don't believe the Universe will correct your simulation if you don't have it just 'right'. If you build an actual _model_ from springs, rubber bands and pencils, the Universe will keep you honest. (O.K., _I_ certainly couldn't build such a model.)

> By running the simulation and adjusting the parameters to get the closest
> possible match of the model's behavior with that of a real arm, we can
> estimate the values of the physical parameters of the real arm. This is the
> basic method of modeling that is behind all the physical sciences and
> engineering. In this world, "simulation" simply means constructing a
> quantitative analog of the system, organized according to a model, and
> running it.

Constructing a _logical structure_ you _hope_ has some relationship to the system. It absolutely is not an analog. That's another serious error of 'traditional' AI.

I didn't mean to imply that simulation was a waste of time. Given the difficulty and cost of constructing and instrumenting models, simulation is often far and away the obvious choice. I only wanted to point out that simulations have no necessary connection with reality, and "results" should be treated accordingly. You learn things from running experiments with real subjects, like the mind reader. Now if you used an analog machine...

Respectfully, Ray Allis

Date: Thu Jun 04, 1992 2:29 pm PST
Subject: ai,speech, reinforcement

[From Rick Marken (920604 14:00)]

eric harnden(900604) says:

> but i feel compelled by my own combination of interests

that's what it's like, being a control system

> to point out that people who engage in AI are not fools.

Nobody said they were. In fact, I've heard nothing but praise for their intellectual skills.

>it is well known that rule-based
>systems can model (or simulate - the usage doesn't affect the issue)
>only a limited subset of what might be termed intelligent behaviors.

I think the PCT position is not that ai type models are claiming to do more than they really can do. Our position is simply that these models are not usually organized as closed loop control systems. No problem there, as long as your goal is not imitation of the behavior of organisms.

> they do, however, provide useful tools
>for information processing, and in fact are alive and well as pre-processors,
>frontends, output analyzers, and other adjuncts to advanced work in
>neural networks.

No doubt. I don't think PCTers are against software engineering. But we believe that if you want to engineer systems that mimic life processes (like intelligent behavior) you will have to take into account some of the "facts of life" -- and one is that the behaviors being imitated are controlled consequences of action. If the term "intelligent behavior" does not refer to controlled consequences of action then current ai models are just fine, both practically and theoretically.

Andy Papanicolaou and Tom Bourbon [920604 13:50] say:

>The hope behind the selection of this specific topic (acquisition
>of, or learning of the skill or the habit of producing correctly
>and reliably new speech sounds) was to elicit specific comments
>that would help us construct a CT model of the process.

This could be a very interesting exercise, especially if you have the tools to do it. I think the first part of this modelling would involve building a vocal tract model -- that produces an acoustic output as does the actual vocal tract. There has been a lot of work in this area so it shouldn't be too hard (theoretically, anyway) to build this. The next part is to decide which variables of this vocal tract model are to be controlled -- or, at least, perceived, so that they can potentially be controlled. I think you will want to control perceived lip positions, tongue positions, larynx opening, etc.-- there are many possible vocal tract variables to perceive and control. Also, of course, you will want to perceive and control aspects of the acoustic results of vocal variations -- things like enery level in the formants, relative energy in formants, spectral width, etc. Again, there are a lot of possibilities but it seems to me like the acoustic linguists know a lot about what might be the important acoustic and vocal variables involved in speech. All you have to do is have a control system controlling each of these variables and determine how these systems should be hierarchically arranged (which determines who's output determines who's reference). This, of course, if the BIG modelling problem. The highest level references in the model will specify what, ultimately, the model is to DO -- perhaps generate phomemes; so the magnitude of each of the highest order references determines the degree to which a phoneme is to be present in perception.

Have fun.

Oded Maler (920604) says:

>The "Reinforcement" in the machine learning context is used in a very
>technical sense without any metaphysical assumption about what "behavior"
>is and all the rest of the S-R psychology you dislike.

Actually, I'll let Bill P. deal with this one. I've got to get outta here. Suffice it to say that it is precisely reinforcement in the "technical sense" that makes no sense as a model of learning or purposive behavior in autonomous agents.

Best regards Rick

Date: Thu Jun 04, 1992 6:26 pm PST
Subject: Re: AI, PCT, and HPCT

(sibun (920604.1600))

[Ray Allis 920604.1200]

Most of the people I know understand that they are SIMULATING intelligence or intelligent behavior (or trying to). After all, the rules of the game say this is to be done with digital computers, which can simulate anything. The problem is that everyone seems to think that means "PRODUCE intelligence". They act as if a simulation is equivalent to the thing simulated.

i think the ``simulate'' vs ``produce'' distinction depends on what one considers ``intelligence.''

But I'd like to point out that I'm not talking about the physical computer, I'm talking about the structure of symbols which the (physical) computer manipulates.

perhaps i don't understand. a simulation is a static arrangement of symbols?

A symbolic structure which is a simulation of a human heart is just as unaffected by reality.

sure, insofar as there is no physical realization of the structure.

but this doesn't seem to be what you're in fact talking about; you say later

The computer operates on an arrangement of symbols, producing another arrangement of symbols.

once you start talking about computers operating on symbol structures, first, you have the symbols physically realized, and second, you are doing something to them--the computer is physically changing a phyisical realization of the structure. so whether a simlation is represented statically in a computer or whether it is in part a result of the process of operating on the representation by the computer (i initially took ``simulation'' to be more of a process than a ``snapshot'' but i guess that's not necessary), it is not unaffected by (physical) reality.

it sounds to me actually as though you are arguing a version of dualism. do you think so?

--penni

Date: Fri Jun 05, 1992 12:08 am PST
Subject: Re: ai,speech, reinforcement

[From Oded Maler 920605]

(Rick Marken 920604):
>Suffice it to say that it is precisely reinforcement in the "technical sense"
>that makes no sense as a model of learning or purposive behavior in
>autonomous agents.

Suppose an agent is in a closed-loop relation with the "world". No matter how it is organized inside, what it perceives of the world results in the "actions" it performs (e.g., muscle contraction) which again, cause certain things in the world, affect the agent's perceptions and so on. "Reinforcement" in this context speaks about how this perception-->action map changes with the history of interaction. If the current map (no matter how realized) fails to achieve the agent's goals (it does not make some high-level reference signal meet their corresponding perceptions, if you like) then it should somehow change.

For the purpose of building machines this is completely legitimate and it has no psychological implications/assumptions more than there are biological ones in "genetic" algorithms. It is just an engineering heuristics that might work or not in certain (real or simulated) situations.

In the context of modeling and HPCT, when you are subscribed to certain organizational principles concerning the perception-->action map, I still think this approach is applicable (maybe if you change the name :-)) to the problem of learning how to servo a complex perceptual variable by changing the connections and parameters in the hierarchy.

The fact that in one discipline/community a certain word carries a baggage of meta-physical assumtions, does not imply that it carries the same meaning in another context.

Best regards --Oded

Date: Fri Jun 05, 1992 6:29 am PST
Subject: Stupid Question

[from Gary Cziko 920605.0900]

Rick Marken (920604 10:00) replied to my 920604.0915:

>> My brain keep saying to me (or at least part of it): "Surely
>>there must be SOME aspect of the cursor which determines the response.

By saying:

>It's the part of your brain that can't think in circles (though you can
>think in circles around me). There IS, indeed, some aspect of the
>cursor that determines the response -- it is the position of the
>cursor. But AT THE SAME TIME the position of the cursor is being determined
>by the response. The significance of this fact is hard to "think through"
>using our usual "lineal" approach to thinking -- a causes b causes c
causes...
>You just have to trust the math on this one.

Looking over my question again, I am now amazed how I could have asked such a question. I've seen and understood the equations. In addition, I've spent a considerable amount of time playing with Bill Powers's Demo2 which shows a "live" control loop in action. It should be perfectly clear to me that "response" is determined by the difference between reference level and perceptual signal (with, of course, cursor also determined by "response"). Yet in spite of all this and my desire to really understand PCT better, I still ask this STUPID question.

This is really simple stuff compared to what the physicists do with quantum mechanics and what the dynamical types do with chaos. And yet people like me who really WANT to understand and aren't particularly dumb still have problems with it.

And while there are more and more people each month grappling with PCT (as shown, for example, by participation in CSGnet), there are probably more and more people each month who are NOT grappling with it and who are basing their work on invalid premises such as perception controls behavior. At times like this I start to wonder if PCT will ever become widely understood.

Maybe somebody out there can cheer me up.--Gary

Date: Fri Jun 05, 1992 7:53 am PST
Subject: Stupid Question -- NOT

[From Rick Marken (920605 08:15)]

Gary Cziko (920605.0900)

That was NOT a stupid qustion. It was precisely the RIGHT question. It is like asking "Why don't all the tree's and people blow off the earth if it's really spinning?". A new model had sure better be able to explain why the "obvious" is not what is happening.

>At times like this I start to wonder if PCT will ever become widely understood.

It will. But maybe not in our lifetime. But it will. How easy could it have been for people to take for grated that the world is turning (at an incredible speed) on its axis and around the sun. Real change isn't easy. I think most people still don't really believe that we evolved from an ape ancestor.

But cheer up. Don't worry about other people understanding PCT. Just enjoy the fact that you understand this precious jewel; desribe it as best you can. Use your understanding and enjoy yourself. There will be a couple more people who understand it each decade or so. Since I got into it, at least one more person in the world learned and understood PCT -- YOU. That is enough to make me happy -- that, and the fact that you are such funny guy too.

The net needs more "stupid" questions like the one you asked.

Best regards Rick

Date: Fri Jun 05, 1992 8:26 am PST
Subject: thinking in circles

[From: Bruce Nevin (Fri 92045 10:53:30)]

(Gary Cziko 920605.0900) --

>Looking over my question again, I am now amazed how I could have asked such
>a question. I've seen and understood the equations. . . .
>Yet in spite of all this and my desire to really understand PCT better, I
>still ask this STUPID question.

>This is really simple stuff compared to what the physicists do with quantum
>mechanics and what the dynamical types do with chaos. And yet people like
>me who really WANT to understand and aren't particularly dumb still have
>problems with it.

>At times like this I start to wonder if PCT will ever become widely
>understood.

I believe the conceptual boggle of "thinking in circles" is analogous to that experienced by people learning to program recursively. I suggest leaning on this analogy when communicating with AI types (or anyone who has learned LISP programming). _The Little LISPer_ illustrates I think a good way to present the fact that there is a conceptual difficulty here and how to get past it. I believe Winston's book on LISP also does. The kind of technical treatment that you see in e.g. Abelson & Sussman's _Structure and Interpretation of Computer Programs_ is less apt because it takes the computer's point of view.

Not to say or imply that I have the knack of thinking in circles yet! But I would point out that this is what the Buddhists mean by mutual causation.

Bruce bn@bbn.com

Date: Fri Jun 05, 1992 9:11 am PST
Subject: Re: reinforcement

[From Rick Marken (920605 10:00)]

Oded Maler (920605) says:

>Suppose an agent is in a closed-loop relation with the "world".
>No matter how it is organized inside, what it perceives of the
>world results in the "actions" it performs (e.g., muscle contraction)

Maybe you're right Gary -- it is impossible.

In a negative feedback closed loop, perceptions are CONTROLLED BY actions, they do NOT cause actions. Actions are "caused" by disturbances to controlled perceptions. Believe it or not.

>which again, cause certain things in the world, affect the agent's
>perceptions and so on.

Which is why what I said above is true.

>"Reinforcement" in this context speaks about
>how this perception-->action map changes with the history of interaction.

Right. But since there is no perception/action map that has anything to do with the behavior of a control system, it sort of obviates the importance of "reinforcement" -- in fact, it suggests that reinforcement (in the sense of changing the probability of a particular output given a particular input) doesn't exist for control systems.

>If the current map (no matter how realized) fails to achieve the agent's
>goals (it does not make some high-level reference signal meet their
>corresponding perceptions, if you like) then it should somehow change.

This works for an open-loop SR system. If this is the kind of system machine leaning types deal with (it is) then reinforcement works just fine. It does NOT work when a system is in a negative feedback situation with respect to it's environment. To see why (experimentally and quantitatively -- no metaphysics necessary) I suggest that you read the papers in chapter 4 of my "Mind Readings" book.

>For the purpose of building machines this is completely legitimate
>and it has no psychological implications/assumptions more than there
>are biological ones in "genetic" algorithms. It is just an engineering
>heuristics that might work or not in certain (real or simulated)
>situations.

Yes, and that heuristic only works for open loop machines -- finite state automata, for example, the most popular machines on the machine learning circuit. Control systems are not finite state automata.

Best regards Rick

Date: Fri Jun 05, 1992 11:14 am PST
From: Dag Forssell / MCI ID: 474-2580
Subject: Direct Mail

[From Dag Forssell (920605 12:15)]

The following is a direct mail letter aimed at the Chief Executive officer of a corporation. The purpose of the letter is to attract and hold the CEO's interest and induce him or her to request further information. This letter is is the very first step, leading to an opportunity to teach PCT in a company and get paid. (I have been preparing full time for 15 months now and have learned a lot).

This rendering is version six. Reactions and suggestions from CSG friends will be appreciated and considered for version seven. I hope some of you will find the ideas useful. How well the letter works remains to be seen, of course. I have already concluded that I did not get far with versions one through four, having mailed 440 letters.

I am anxious to be correct in what I say. The response rate will be low, but those who do reply will know what they are getting into. When I am teaching and get attacked by the company's people experts (which is bound to happen), I am prepared, and so is the CEO.

This is a closed loop feedback learning process, with much action sending mail and little feedback from replies. So you reorganize (rewrite), and try again.

Page 1

(Purposeful LeadershipTM letterhead)

Adam Smith, CEO June 5, 1992
Smith & Smith Inc.
1000 Main Street
Smithsville, State, Zip

Dear Mr. Smith:

I am writing to introduce you personally to a new perspective on human interactions, which has many implications for leadership. This perspective gives an executive insight that allows him or her to inform, influence, align and lead people with mutual respect. He or she can teach people to be more effective and more cooperative. Employees can be more effective and satisfied, while the company as a whole responds better to the leader's direction and becomes more productive.

This perspective also will make it much easier to understand and teach

Total Quality Management programs, such as the Deming Management Philosophy.

Describing this perspective so you get the point immediately is a Catch-22 challenge, because it is a different concept altogether from what predominates in our world today. The language we use reflects the currently predominant concept. To describe something fundamentally different with that language misses the mark easily because the words mean something conceptually different to the reader than they do to the writer.

Let me use an illustrative analogy:

In an era when "everyone knew" that the earth was flat, scientific explanations were developed for navigation and astronomy. Many problems with these explanations persisted, but people worked around them. There was no alternative. The explanations were taught to succeeding generations by experts, who derived status from their knowledge. Non-experts took it all for granted without much thought.

I cannot say what "everyone knows" about human behavior, but experts on the subject employ a 17th century perspective of cause and effect to guide their research. Any book on experimental psychology tells you that the way to learn about behavior is to set up an experiment, then vary the stimulus (independent variable) and watch the response (dependent variable.) With this scientific method our experts have made many experiments and formulated many explanations which have found their way into our language, culture and management practices. Non-experts take these explanations for granted without much thought.

Many problems with these explanations persist despite all the research, but people work around them. There is no alternative. Our understanding of what motivates people is very poor. We clearly lack a good model or "paradigm" to help us understand why people do what they do. In our ignorance, we tend to spend our energies in debilitating conflict again and again, instead of in productive cooperation.

Page 2

When Copernicus and then Galileo introduced the new fundamental insight that the earth is round (it had been round all along), the problems of navigation and astronomy were placed in a bright new light. The new insight did not invalidate the common sense observation that the earth appears flat locally, but science moved from a dead end to progress, which in a few centuries has brought us far.

But the experts of the day could not (and did not want to) comprehend the new paradigm, because they had already internalized the flat paradigm in all its details as their personal reality. With time, the experts died off, and new ones grew up, embracing the new paradigm on its merits because it solved many of those persistent problems. They internalized the new perspective, and science progressed from there.

Isaac Newton's "Principia Mathematica," published fifty years after Galileo, was resisted also for similar reasons. It took fifty years for it to be fully accepted. Looking back, we take it for granted. (As a result we are able to predictably go to the moon)! The evolution of science is much more than a steady accumulation of knowledge!1

1Footnote: The phenomenon and process is described in Thomas Kuhn's seminal book: "The Structure of Scientific Revolutions," which introduced the term "paradigm."

The 20th century understanding of the phenomenon of control, includes the appearance of simple cause-effect relationships in a larger framework. Applied to people (people have been controlling all along) and developed into a detailed model by William T. Powers, it provides a fundamental insight that puts the problems that result from human interactions in a bright new light. Because the new insight is basic, it has consequences in many ways in many areas of an organization and the results from using this insight can be dramatic and far reaching.

There have always been natural leaders, successful salesmen, wise parents and good communicators. But it is rare that they can explain what they do and why. Their insight and skill is intuitive. The benefit of a good understanding of control is that you gain explicit clarity and can learn to function as well as the intuitively wise people. With practice even better, since you will know what you are doing.

The new insight into the phenomenon of control does not invalidate any wise common sense observation or practice. It just provides an enhanced understanding of seemingly intractable problems. It provides new diagnostic tools and shows why cookbook formulas for behavior are inappropriate. It requires that you think for yourself. You (and other non-specialists) can evaluate it better than established "experts" because their existing education and status will tend to blind them to the new perspective. (In his book: "Future Edge", Joel Barker says: "When a paradigm shifts, everyone goes back to zero." The old knowledge is obsolete. The experts lose their advantage).

This perspective on control is already well developed. But no doubt it will take time - into the 21st century - before this successful development is embraced by a majority of experts. You can take advantage of what "everyone will know" in the 21st century right now to improve your company's competitive position. But to do it, you must be willing to think for yourself. You will actively participate in a scientific revolution.

The Purposeful LeadershipTM programs explain and translate this new perspective into skillful use of diagnostic tools that give you the capability to work on productivity. That includes effective communication, teaching effectiveness, resolving conflict, supporting self-motivation in employees, team building, Total Quality Management, leadership insights, effective performance appraisals, effective selling concepts, and development of corporate and individual mission statements.

Page 3

The basic principles can be taught in a day to any attentive person, who can also verify them. People trained in the "hard" sciences will appreciate the scientific approach and elegant simplicity of the program, and everyone will be able to begin applying the principles as soon as they understand the underlying model of control and have had some instruction and practice with applications.

Some people will think that "understanding the phenomenon of control" promises a new way to control other people. It is precisely the other way around. We show how people control themselves at all times. When you understand control you can work with people, rather than get into conflict despite the best of intentions.

I have 25 years experience in engineering, manufacturing, financial and marketing management. My formal education includes an MBA from the University of Southern California and a Masters degree in Mechanical Engineering from Sweden.

I will be pleased to send you an introductory 39 minute audio tape (with script for reading and illustrations) which explains the background, scope and applications of our programs. The introduction includes a demonstration/test which will allow you to determine if your associates can recognize control in action.

When you receive the introduction, I think you will find the demonstration both enlightening and entertaining. Please feel free to share it with your technical, operations and sales managers at any level for their evaluation. This is a win/win program to greatly increase the understanding and effectiveness of anyone who deals with people.

Sincerely,

Dag Forsselll

(Simple form for request for tape & script).

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

Date: Fri Jun 05, 1992 2:55 pm PST
Subject: Re: AI, PCT, and HPCT

[Ray Allis 920605.1230]

> (sibun (920604.1600))

You're really gonna make me work at this, aren't you? :-)

> A symbolic structure which is a simulation of a
> human heart is just as unaffected by reality.
>

> sure, insofar as there is no physical realization of the structure.
>
> but this doesn't seem to be what you're in fact talking about; you say later

Yep, that is what I'm talking about. The idea is that deduction is
only concerned with the FORM of a structure (argument); the content
is not relevant. (Even deliberately removed as in algebra.)

>   The computer operates on an arrangement of
>   symbols, producing another arrangement of symbols.
>
> once you start talking about computers operating on symbol structures,
> first, you have the symbols physically realized, and second, you are
> doing something to them--the computer is physically changing a
> phyisical realization of the structure. so whether a simlation is
> represented statically in a computer or whether it is in part a result
> of the process of operating on the representation by the computer (i
> initially took ``simulation'' to be more of a process than a
> ``snapshot'' but i guess that's not necessary), it is not unaffected
> by (physical) reality.
>

Correct. But it doesn't matter. A (digital) computer running a
program is affected by (physical) reality. The (physical instances of)
symbols (as packets of electrons or whatever) are real and are affected
by (physical) reality. But it's some electrons which are affected, not
whatever it was that was symbolized by them. That relationship only
exists in your mind. The computer pushes electrons, not quantities.
1 + 1 = 2 is a statement in formal logic. Formal logic is content-free.
The symbols are only there so you can have interrelationships
among them. They don't even have to symbolize anything (i.e., be
symbols at all). The FORM's the thing.

We owe this sort of blind spot to the ancient Greeks. Maybe it's fair
to use something as old as a syllogism to illustrate. All men are
mortal. Socrates is a man. Therefore Socrates is mortal. Great, hardly
anyone will disagree. Now suppose I say - All women are irrational.
Penni is a woman. Therefore Penni is irrational. Does that give you
any problem? If so, is your problem with the form of the argument?

No, (I bet) the problem is that you perceive a conflict between your
experience and the MEANING of the first premise. Note that the meaning
is not a property of the set of symbols, but rather exists only in your
mind. Formal logic will not help you here. And formal logic is
what a computer simulation is.

Models, as opposed to digital computer simulations, can provide new
experience. It happens that if you pour one liter of alcohol and one
liter of water into a two liter container, you discover that you don't
quite have two liters of mixture. Hmmm. This is not discoverable by a
digital computer simulation. (Of course you can account for it once
you know.)

> it sounds to me actually as though you are arguing a version of
> dualism. do you think so?

Gee, I hope not! I think my mind is a result of the operation of my central nervous system. Even-numbered days I believe that may include my whole body.

Ray Allis

Date: Sat Jun 06, 1992 4:22 am PST
Subject: AGARD Presentation

[Martin Taylor 920606 0815]

Well, I'm back, but I haven't had time to read the 400-odd messages that piled up over the last six weeks. I want to send this out now, even though I probably won't respond to comments on it until I have read the older posting.

As some of you know, the main (official) reason for my extended absence was to present a paper at an AGARD workshop in Madrid on "Advanced Aircraft Interfaces."
Bill Powers made a nice comment about the first version of the written paper, and I gather that one or two of you asked for copies while I was away. The main theme of the paper as written was to introduce Perceptual Control Theory and the Layered Protocol theory of communication (LP) and to use them to indicate when and how to use voice communication in an aircraft cockpit. But that turned out not to be what I actually presented. I did introduce PCT and LP, much as written, but ignored the voice in favour of something that seems more interesting: the graceful use of automated functions in the cockpit. On the first day or two of the workshop, several talkers made the point that pilots had a difficult time accepting automated functions beyond the most simple, although they indicated in questionnaires that they wanted them. What they did not want was for the automated functions to take decisions that the pilots would rather take for themselves at critical moments, although the automated function could perform non-critical duties.

It struck me as I listened to these presentations that the PCT view provided a framework within which this problem was perhaps soluble, so that was what I discussed instead of the voice.

Both the plane and the pilot are conceived as hierarchic control systems, the plane's upper-level references being set either by the designer or by the pilot. A system like an autopilot has a reference to keep the plane on a certain heading at a certain altitude and with a stable attitude regardless of winds. The pilot resets this reference from time to time, or it could be reset from a higher-level sequence controller that alters the desired heading and altitude whenever the plane reaches a waypoint. In this case, the pilot sets up the waypoint sequence in geographic coordinates, thus providing the reference sequence to the higher -level ECS of the plane. In both cases, if we think of the pilot and plane as one single hierarchic control system, the plane's chunk simply takes over the function of performing the task of satisfying the references provided by the pilot. Indeed, if the pilot trusts the autopilot, there is no need even for her to perceive where the plane is going, so that her "control" is non-existent. He is in the situation where her environment is sufficiently stable that she can operate (at that low level) without feedback. The plane WILL go where she asked, and she does not have to worry about checking how it is progressing. (For the possible effects of working that way, remember the Korean airliner that was shot down by the Soviets, in part, the investigators claimed, because the pilot entered a wrong course or waypoint for the autopilot, and trusted it to be right.)

If the pilot has delegated control, the plane taking over the particular function, the pilot tends to lose "situation awareness." He could control the function if he wanted to, but since he is not, neither is he acquiring the sensory information that would allow him to get the perception that he could be controlling. He does not perceive what is going on. The pilot's re-acquisition of situation awareness when retaking control of an automated function is a significant problem. The autopilot is switched out of the control loop, and the pilot's own lower-level control systems take over the maintenance of heading, altitude, and attitude. One significant reason for this to happen is collision avoidance, where situation awareness is critical.

The view of the automated function being switched in or out of the loop in alternation with the equivalent part of the pilot's hierarchy is almost inevitable with conventional approaches to the problem. But PCT offers a different solution (clued by Rick Marken's observations in April that two competing control systems can provide a lower variance than either working alone). Imagine that instead of a simple switch that sends a reference signal either to part of the plane's hierarchy or to part of the pilot's, the reference signal is sent always to both. If the pilot is choosing not to control, the gain in her part of the loop is zero (and as in the other case she may not even be aquiring the appropriate sensory information). The gain in the aircraft's part of the loop is adequate to maintain course against external disturbances.

But it is possible for the pilot to set his gain to some low value other than zero, and "shadow" the aircraft's control. The aircraft could sense this in two ways. One way is that the pilot's attempts to control would set up a conflict in the lower-level systems that actually drive the plane's control surfaces. The result would be a persistent failure to achieve the deisred percept, if the pilot's references differed from those of the plane. The second is that in contrast to ordinary disturbances, the pilot's actions can be directly sensed by the plane and the plane can act as the PCT experimenter that we often discuss, TESTing whether and to what degree the pilot is controlling. So long as the pilot's gain remains low, the automated system would keep its own gain high, but as soon as the pilot's gain increased (indicated his "insistence" [1] that he take control), the plane would drop the gain of the automated system, perhaps to zero. The pilot is, at low gain, maintaining situation awareness, or regaining it preparatory to taking control.

There is a continuum here, as the plane's gain decreases, between the plane performing the function, assisting (and perhaps training) the pilot to perform it, and getting out of the way to let the pilot do what she wants. There is no need for the pilot to switch automated functions in and out; they are in by default, but as soon as the pilot starts controlling what they control [2], they gracefully get out of the way. What the pilot can switch in or out, or alter in a continuous way, is the sensitivity of the plane to the pilot's insistence on control. Thus a novice pilot could set a high level, asking the plane to do what it thinks proper even though she requires it moderately strongly to do something else, whereas an expert would want it to get out of the way as soon as she started controlling.

Shifts of control locus need affect only a small part of the hierarchy. The pilot's choice to control the course of the plane does not indicate that he must control the positions of individual control surfaces. Indeed, in modern high-performance planes, any attempt by the pilot to do so would lead to disaster very quickly. Only the course control would shift between plane and pilot, leaving the plane in control of the actual movements of the surfaces. And the TEST allows the plane to know at what level the pilot does desire to take control (maybe not quickly enough to get it right immediately, so it might have to relinquish control over several levels, regaining it by default over those levels the pilot fails to control).

There seemed to be a positive reaction from one or two participants at the meeting, and a demand for the paper. Maybe something will come of it.

----
[1] In the talk I used the term "insistence" as a generalization of "gain", because it seems to me to be more appropriate for control loops that contain categorical boundaries, and to be an adequate term for continuous control loops. The pilot is "insistent" that the target airport is Bologna rather than Roma, but it is hard to justify the term "gain" for that kind of control. Unless there is contrary comment, I think I will continue to use "insistence" for generalized "gain" in discussions in this group.

[2] There's a problem with cumbersome use of language here. Of course the plane and the pilot cannot be controlling the same thing, because each has a "personal" percept, which is actually being controlled. But to a large extent each of the two percepts depends on the same environmental complex, and it is much easier to say that plane and pilot are controlling the same environmental complex (the course, for example) than to say something like "the pilot starts controlling for a percept affected by an environmental variable very close to the complex that affects the percept for which the plane is controlling." But the easy form of language can lead one into sloppy thinking, as is often the case with language.

Martin -- I'd rather be still in Spain.

Date: Mon Jun 08, 1992 6:40 am PST
Subject: Modeling & simulation

[From Bill Powers (920608.0100)]

My host computer went off the air last Friday just as I was about to send a message saying that I'd be gone for a couple of days. I'm back.

Chuck Tucker asked me for a copy of my observations about Coach in high school. I tried to find them and couldn't (I searched through four months of posts and found only that I talk too much). If anyone has a better filing system than I do (easy to accomplish), would it send a copy to Chuck at n050024@univscm.bitnet? (Those are zeros, not ohs).

J. Marvin Brown (CT linguist from American School in Thailand) will be visiting here later today and tomorrow. Messages from other linguists (or anyone else) welcome. Marvin is not on the net.

Ray Allis (920604.1200) --
Penni Sibun (920604.1600) --

I think there are four universes of discourse overlapping here (at least) in the discussion of modeling, simulation, etc.

Ray is arguing that symbols manipulated in a computer have nothing to do with reality, in that symbol-manipulations can contain no surprises and the rules for manipulation contain all possible deductions from them. Penni is pointing out that the computer itself is a physical reality and that the symbols themselves are quite real as is the structure (in the computer) that is realized in terms of symbols. That's on the one hand.

On the other hand, in the world of CT modeling so far, the computer is not used to manipulate symbols at all: it's being used as a numerical analog computer dealing with variable quantities. Furthermore, I haven't yet pointed out that in many of our simulations one component is a real person performing real behaviors (moving a control handle and perceiving one or more moving targets and cursors, or other types of displays) in a continuous manner while the remainder of the setup runs as a program inside the computer. So the program has no advance knowledge of what the person is going to do next, nor does it know how the person will interpret and react to the display. The simulation aspect then amounts to trying to write another program which, when substituted for the real person in the same experiment, will produce the same record of behaviors that the real person produced (or will produce in the future, under new experimental conditions).

In the case of CT experiments, I object to classifying the processes inside the computer as logical deductions because the same processes could be carried out by resistors and capacitors and operational amplifiers in a purely analog setup (or any other means of analog computing such as pneumatic or mechanical).

On this same hand, Rick Marken has said that a link between the computer simulation and reality is continually present through the experimenter, who is comparing what the computer does with what a real person does in the same experimental situation, with a very critical eye for differences. The simulation does not run free of influences from reality; in fact the simulation is set up to preserve critical correspondences with the external world or the real system. The better simulations attempt to reproduce many correspondences, not just those at input and output.

On the third hand, there is the fact that either in a digital computer or in an analog computer, a human being must be present to interpret the inputs and the results. Digital and electronic analog computers don't actually manipulate any symbols or perform any logic. They manipulate voltages according to properties of electronic circuits. These voltages have no meanings, because neither kind of computer has any experiences (except other voltages) to become the meanings of the voltages.

The programmer of either kind of computer assigns meanings to arbitrary voltages used as symbols, these meanings coming from the programmer's experiences, not the computer's. The computer then manipulates the voltages according to the physical properties of its components and its programming, and produces, as output, more voltages. The programmer then looks up the assigned meanings of the output voltages (preassigned by the programmer) and interprets the results accordingly.

In digital computing, the inputs and outputs are usually cleverly programmed so that the assigned meanings of the voltages are entered and read out as codes displayed to look like words or numbers that remind the programmer of their intended meanings. These words or numbers have no meaning in the computer, but they do have meaning for the programmer: the order in which different words or numbers come out of the computer and the words or numbers themselves suggest new thoughts and experiences to the programmer (or, of course, the user).

In analog computing, the input and output voltages (now continuously variable instead of on-off) are assigned meanings as physical variables: the inputs might be said to correspond, for example, to temperatures, pressures, and densities, while the output voltages might be said to correspond to rates of heat flow in various places. The inputs are set up, the computer is set in motion, and a tracing of the output voltages is delivered. The output voltages are plotted against time with the coordinates labelled as joules per second, heat exchanger, and so on, which is the meaning of the output to the human programmer, but not to the computer. The assignment of meaning by the human programmer is more obvious in an analog computer, which does not produce outputs in the form of recognizeable words or numbers.

On the fourth hand, there is the question of the nature of the human programmer. If we accept the idea that the programmer's brain is a neural machine, as I do, and that all aspects of experience and action are embodied in and by this machine, then it follows that both symbols and meanings exist in this machine as neural signals, trains of impulses occurring at variable frequencies. Reality then becomes a representation inside this machine, knowable directly to the machine, but corresponding to something not directly knowable by the machine.

The problem of modeling and simulation then becomes one of creating an analog of the brain-machine in relationship to a hypothetical physical reality outside the brain-machine. This hypothetical reality also exists in the brain, in the form of neural signals being handled by neural functions that exist as synaptic connections and chemical reactions. The brain sometimes attempts to mirror itself in artificial machines, giving them internal operations that make sense to the brain. It naturally constructs these artefacts to accept inputs and produce outputs that can be perceived by the brain, that perform operations recognizeable by the brain, that remind the brain of experiences other than those inputs and outputs. So the brain, which operates in terms of basically identical signals handled by processes it cannot perceive directly, tries to discover its own internal organization by creating external artefacts in which those invisible operations can be made visible -- and which also operate in terms of basically identical signals.

There is, finally, a fifth hand, but of a different kind. All these conceptions of a brain, including that of a brain as a neuronal machine, exist in the consciousness of an observer-manipulator which is ourselves. Before we learn anything formal about the physical world or machines, the experienced world exists just as it exists. Whatever we guess to exist behind this world of direct experience, it is always direct experience to which we return in evaluating any guess. Any conjecture that directly denies direct experience is simply wrong. If you want to say that some aspect of direct experience is imaginary, for example, you must then explain how it is that this particular experience was imagined: if you conclude that it was not imagined exactly as it appeared, then something is wrong with your reasoning. No person can do other than rely on direct experience as the final arbiter. If it seems to you that someone else is mistaken, then that is the case: that is, in fact, how it seems to you. If you think the other person is right, then that, too, is how it seems to you. There is simply no way to deny direct experience. If I think therefore I am, (or therefore I am not) then that is what I did in fact think. It does not matter whether, in some hypothetical objective world independent of my experience, I really am. Such objective facts and their implications are conclusions which we test, ultimately, against experience.

It's not easy to keep straight where we stand in the midst of these different universes of discourse. I think that we're trying to achieve consistency among them, so that moving from one to another doesn't create contradictions among them. No one of these points of view is so important that we can afford to ignore the others; from each point of view we see a different facet of the mystery of being conscious in a here and now.

Best to all, Bill P.

Date: Mon Jun 08, 1992 8:10 am PST
Subject: A new control method

Hi :

I am a new group member. In this letter, A new control method which is similar to the inverse dynamic (a nonlinear control method) will be described simply. If anybody has touched or researched this method, or has other different solutions, please give me the suggestions. I just want to find friends.

1. A brief primal concepts :

Given a linear discrete-time system model in the z-domain type

```
   Y(z)      b0 + b1/z + b2/(z^2) + ... + bm/(z^m)
  ------ = ---------------------------------------
   U(z)      1 + a1/z + a2/(z^2) + ... + an/(z^n)
```

Where y() is the system output and u() is the control signal.
At some time k, The time-series relation between outputs and control is

```
   y(k) + a1*y(k-1) + a2*y(k-2) + ... + an*y(k-n) =
          b0*u(k) + b1*u(k-1) + b2*u(k-2) + ... + bm*u(k-m)
```

i.e.

```
   y(k) = b0*u(k) + b1*u(k-1) + ... + bm*u(k-m) -
                    a1*y(k-1) - ... - an*y(k-n)               (1)
```

If this time is k, then we can forecast the future outputs y(k+1), y(k+2) , ... through equation (1),

```
   y(k+1) = b0*u(k+1) + b1*u(k) + ... + bm*u(k-m+1) -
            a1*y(k) - a2*y(k-1) - ... -an*y(k-n+1)
   y(k+2) = b0*u(k+2) + b1*u(k+1) + ... + bm*u(k-m+2) -
            a1*y(k+1) - a2*y(k) - ... - an*y(k-n+2)
             .
             .
             .
   y(k+i) = b0*u(k+i) + b1*u(k+i-1) + ... + bm*u(k-m+i) -
            a1*y(k+i-1) - a2*y(k+i-2) - ... - an*y(k-n+i)
```

Where y(k+1), y(k+2), ..., y(k+i) and u(k+1), u(k+2), ..., u(k+i) are all unknown, move the unknown outputs to the left by recursive method,

```
   y(k+1) = b0*u(k+1) + b1*u(k) + ... + bm*u(k-m+1) -
            a1*y(k) - a2*y(k-1) - ... -an*y(k-n+1)
   y(k+2) = b0*u(k+2) + (b1-a1*b0)*u(k+1) + (b2-a1*b1)*u(k) + ...
            + bm*u(k-m+1) + (a1*a1-a2)*y(k) + (a1*a2-a3)*y(k-1) +
            ... - a1*an*y(k-n+1)
         .
         .
         .
   y(k+i) = b0*u(k+i) + b1'*u(k+i-1) + ...... +
            a1'*y(k) + .... + an'*y(k-n+1)
```

Where u(k), u(k-1), ..., u(k-m+1), and y(k), y(k-1), ..., y(k-n+1) are known because this time is k, so

```
   y(k+1) = b0*u(k+1) + c1
   y(k+2) = b0*u(k+2) + (b1-a1*b0)*u(k+1) + c2
         .
         .
   y(k+i) = b0*u(k+i) + b1'*u(k+i-1) + ... + b(i-1)'* u(k+1) + ci
```

c1, c2, ..., ci are known, arranging these euqation in matrices

```
   / y(k+1) \        / u(k+1) \
   | y(k+2) |        | u(k+2) |
   |   .    | =  T * |   .    | + C
   |   .    |        |   .    |
   \ y(k+i) /        \ u(k+i) /
```

Where T is an i by i coefficient matrix, C is equal to [c1,c2,..,ci]'. If n' <= i and n' > 0, the following equation can be picked out

$$\begin{pmatrix} y(k+i-n') \\ y(k+i-n'+1) \\ . \\ . \\ y(k+i) \end{pmatrix} = T' * \begin{pmatrix} u(k+1) \\ u(k+2) \\ . \\ . \\ u(k+i) \end{pmatrix} + C' \tag{2}$$

T' is the submatrix of T and (n'+1) by i, C' is the subvector of C and its length is (n'+1).

Consider the basic step-response control, the system outputs are always desired to approximate a fixed reference. On the steady state, the system outputs are equal to the reference and the control signals are usually kept at a fixed value. Realizing the relation between the future system outputs and the control signals, we can obtain an optimal control signal to make the future outputs approximate the reference (stable) and the control signal approximate a fixed value (finite energy).

A error function is defined as follows to get the optimal control signal. (Let the refrence be r(k) at time k.)

$$E = \begin{pmatrix} y(k+i-n') \\ y(k+i-n'+1) \\ . \\ . \\ y(k+i) \\ u(k+i-m') \\ u(k+i-m'+1) \\ . \\ . \\ u(k+i) \end{pmatrix} - \begin{pmatrix} r(k) \\ r(k) \\ . \\ . \\ r(k) \\ us \\ us \\ . \\ . \\ us \end{pmatrix} \qquad us = \frac{1+a1+...+an}{b0+b1+...+bm} * r(k) \tag{3}$$

and

$$n' = \begin{cases} i-1 & \text{if } i <= n \\ n & \text{if } i > n \end{cases} \qquad m' = \begin{cases} i-1 & \text{if } i <= m \\ m & \text{if } i > m \end{cases}$$

Substitute equation (2) into (3), the relation between the error function and the control signal is

$$E = \begin{pmatrix} T' \\ 0,\ I \end{pmatrix} * \begin{pmatrix} u(k+1) \\ u(k+2) \\ . \\ . \\ . \\ u(k+i) \end{pmatrix} + \left(\begin{pmatrix} C \\ 0 \end{pmatrix} - \begin{pmatrix} r(k) \\ r(k) \\ . \\ . \\ r(k) \\ us \\ us \\ . \\ . \\ us \end{pmatrix} \right) \tag{4}$$

Where I is an unit matrix, and 0's are zero matrices.
Minimize (E'*E) by the least-square method (QR decomposition as usual), the corresponding control signal u(k+1), u(k+2), ..., u(k+i) is the optimal values which makes future outputs y(k+1), y(k+2), ..., y(k+i) approximate r(k) and the control signal

approximate a fixed value us as possible. However, only u(k+1) is selected to apply to the controlled system at time k+1, and the same approaches is reviewed at the following time to get the optimal u(k+2).

In equation (4), variable i is called as CONTROLLER ORDER, which is suggested to be limited from 1 to (n+m+1). When i=(n+m+1), the minimum (E'*E) is zero and the stability of the controlled system is sure.

Besides, a parameter w is jointed into the error function to improve the output performance, which is called WEIGHT.

$$E = \begin{pmatrix} w*y(k+i-n') \\ w*y(k+i-n'+1) \\ . \\ . \\ w*y(k+i) \\ u(k+i-m') \\ u(k+i-m'+1) \\ . \\ . \\ u(k+i) \end{pmatrix} - \begin{pmatrix} w*r(k) \\ w*r(k) \\ . \\ . \\ w*r(k) \\ us \\ us \\ . \\ . \\ us \end{pmatrix}$$

The weight is usually greater than one except some nonminimum-phase systems. When i=1 and w=infinity, this method is equal to the inverse dynamic.

2. Remarks

1. The concepts of this method is based on SYSTEM FORECAST. The most of human control behaviors is also based on the similar mode, and we call that 'EXPERIENCE'.

2. This method provides very rapid analyses for any control system, whether stable or unstable, minimumum-phase or nonminimum-phase.

3. The robustness is worse under model noise, an improvement is to select new control signal du, du is defined as follows :

   $$du(k) = u(k) - u(k-1)$$

   in the z-domain type

   $$dU(z) = U(z) - U(z)/z$$

   Through the new control signal, the steady-state error due to model noise is deleted.

4. An interesting skill : for varying references, if the future references are known at every sample, the error function can be modified to

$$E = \begin{pmatrix} y(k+i-n') \\ y(k+i-n'+1) \\ . \\ . \\ y(k+i) \\ u(k+i-m') \\ u(k+i-m'+1) \\ . \\ . \\ u(k+i) \end{pmatrix} - \begin{pmatrix} r(k+i-n') \\ r(k+i-n'+1) \\ . \\ . \\ r(k+i) \\ us(k+i-m') \\ us(k+i-m'+1) \\ . \\ . \\ us(k+i) \end{pmatrix} \qquad us(j) = \frac{1+a1+...+an}{b0+b1+...+bm} * r(j)$$

This skill results in a null-phase-delay response.

5. This control method is easy to link an on-line system-identification implement, an adaptive control system can be built up to retrieve the worse robustness.

6. For multi-input-multi-output systems, a very enormous matrix can be foresight, and the null diagonal elements may be existed in the QR decomposition.

Bill Chen (Kuo-Feng Chen) of NCTU in Taiwan.

Date: Mon Jun 08, 1992 8:21 am PST
Subject: Re: standards

[Martin Taylor 920607 1710]

On reading through my 400-message backlog of mail, I find several messages on standards and how they should be interpreted in the PCT world. Having got as far as Marken (920520 19:00), I find no-one has interpreted standards as I would--and I am surprised Bruce Nevin hasn't done so, because my reading of standards is rather like his on the external nature of language.

In interpersonal communication of any kind, including language, one can best achieve control of one's percepts if one has some notion of what the other is likely to do that affects your sensory organs. If you don't want to perceive yourself being hit with a 2x4, you don't antagonize a Hell's Angel. You model the partner in some way. It seems to me that standards allow you to pre-empt a possibly painful random reorganization by permitting you to set references that are appropriate if the other behaves in a conventionalized way--according to standards. Likewise, if you behave according to standards, your references will be set so that your observable behaviour conforms to the expectations of the other--they will know what you are controlling for at the relevant level, and will be able to interpret low-level acts/behaviours as supporting that control.

If there are any absolute standards, they will be those that have allowed the social groups using them to survive and prosper. A standard that allowed group members to kill one another for fun is not one that is likely to be found in a long-surviving group. Our standards have been evolving since at least the time humanoids diverged from other primates, and there are clearly some sets of standards that work well together but are different from other sets that also work well together. One standard that worked well when relatively isolated tribes wandered around competing for resources involved wariness and intolerance for people not of one's own group. Killing them meant more for one's own group. Racism comes from this. But recently there has come to be only one communicating group in the world, and this long-useful standard seems to be one that will not allow this single group to survive long if it maintains its currency as a model for how to set a reference level.

Standards for grammatical usage seem to have exactly the same theoretical standing as standards for good social behaviour. One sets references for using correct grammar because it eases the task of communicating partners who use the same standards. If a subgroup uses different standards, there's no problem except that their communication with the main group becomes less effective. If one person decides on a different set of reference levels (such as not capitalizing the initial letters of sentences), they cause communication problems with all their partners. There's no moral good or bad about it, only a consideration of efficiency.

We can't do without standards in a time-limited social world.

Martin

Date: Mon Jun 08, 1992 10:53 am PST
Subject: Pengi & HPCT

> "The-block-I'm-pushing
> The-corridor-I'm-running-along
> ...
> The-bee-that-is-heading-along-the-wall-that-I'm-on-the-other-side-of."

> It seems to me that being able to recognize such things entails a very complex perceptual system, capable of discriminating, recognizing, and naming objects (block), processes (pushing), agency (I'm pushing), relationships (I pushing block), and so on.

No, actually, the intention is that these are theorist's names for things that are detected by the agent in much simpler ways, as you suggest. They are not intended to be treated compositionally; ergo being able to detect The-block-I'm-pushing does not entail being able to detect pushing. This point is probably brought out more clearly in my thesis.

> A couple of years ago I wrote a program for Clark McPhail (sociology) at U of IL that simulates the movement of actors through a field of obstacles and other actors.

This sounds like a bunch of current work in the artificial life movement. If you are looking for a place to publish, their annual conference might be a good place to send something.

Date: Mon Jun 08, 1992 11:10 am PST
Subject: modeling

[From Rick Marken (920606)]

Well, I've just got to get into this discussion of modeling that's going on between penni s., Ray Allis and Bill P.

I guess I don't understand Penni's or Ray's position on modeling (or simulation or whatever). In am having a particularly tough time understanding Ray's position -- I must not understand because I think he is saying that the best way to understand phenomena is to build them (that's what I am picking up from his definition of models). For example, in Ray's post of

920605.1230 he says:

>Models, as opposed to digital computer simulations, can provide new
>experience. It happens that if you pour one liter of alcohol and one
>liter of water into a two liter container, you discover that you don't
>quite have two liters of mixture. Hmmm. This is not discoverable by a
>digital computer simulation. (Of course you can account for it once
>you know.)

Perhaps there is an expression problem here but it sounds like "pouring alcohol into water" is being proposed as a model of alcohol-water mixture. I can't believe that science would have gotten very far with this approach to modeling. Maybe Ray could explain what the above statement is supposed to mean. Does it mean that with a good computer simulation I cannot predict phenomena that I have not yet observed? This must be wrong since models like relativity were predicting new phenomena all the time (like the light bending near large masses) and relativity can be simulated on a computer; it is difficult to build what I am understanding as your proposed type of model of a relativistic cosmos.

Apparently, your version of a model does not have to be built of the material that we know the actual system that exibits the phenomenon to be built of; we can "simulate" components to some extent. For example, organisms, which exhibit the phenomenon of control can apparently be built from rubber bands and tubing rather than actual muscle cells and veins. I understand this about your version of models from the following:

[Ray Allis 920604.1200] (in reply to a comment by Bill P.)

> I don't believe the Universe will
>correct your simulation if you don't have it just 'right'. If you build
>an actual _model_ from springs, rubber bands and pencils, the Universe
>will keep you honest. (O.K., _I_ certainly couldn't build such a model.)

This conflicts somewhat with my idea of how science works. It seems to me that we begin by making observations and then invent models made of unseen entities and functions that can produce these phenomena. If it's a good model (regardless of what it is made of) it will produce a quatitative match to the observed phenomenon. The model also predicts phenomena that we have not yet seen (just as the relativity equations predicted the light bending). We then test these predictions. It is here that the universe keeps us honest-- if we observe the phenomenon as expected, then we gain a bit more confidence in the model. If not, we change the model as necessary. The scientist keeps the model honest -- mother nature keeps the phenomena honest. If light had not bent near a mass, then the model would have to be changed. The universe doesn't care whether your model is correct or not, only the scientist cares.

I'm getting the impression that Ray (and perhaps penni also) is criticising ai for not building its models from physical analogs of the stuff that carries out intelligent behavior. I think this misses the point somewhat. It's like arguing that a robot model would be better if it were made out of flesh and blood. Now that I think of it, I seem to recall philosophers (like Searle) criticizing ai by saying things like "ai isn't a good model of intelligence because computers don't have the vegetative requirements of humans". I (mildly) disagree.

The problem with ai (I think) is that it leaves out models of the environment in which "intelligent behavior" is carried out. All we know of the "reality" in which intelligent systems carry out their behavior is a model anyway. We use Newton's model of physical reality in our control simulations, usually. There is no need to actually build intelligent systems out of "real" stuff; it would't be feasible anyway: and since it is not, how would one know what to leave out? or what to substitute (are rubber bands really the right substitute for muscles? do they fatigue in the same way?).

AI models often ignore the complete, relevent situation in which they are suposed to behave. For example, ai models assume that outputs go into a non-physical, non-time-dependent vacuum that changes the "input" just like that; what they are leaving out is what physicists tell us is a world of forces, inertias, momenta, etc; that is, they are leaving out the extremely successful models we already have of what has been referred to in this discussion as "reality" (these physical models are, of course, actually models of the causes of our perceptual experience).

I think there is much to be learned from models of intelligent systems that are built out of physical components (the kind that Ray seems to be advocating). But such models are still simulations (according to Ray's definition). For example, you are using the elastic properties of rubber to simulate the elastic properties of muscle (just as you could use a differential equation or a computer program to model these elastic properties -- the relevent part of the model is the functional relationship between changing variables -- function which can be implemented quite nicely in a computer).

Just for sheer power and non-messiness I have to come out strongly in favor of computer modeling as an approach to understanding the phenomenon of purposeful behavior in organisms. Just because ai people have done an incomplete job of such modelling does not mean that something is wrong with that approach to modeling itself. And I think many of the ai algorithms (such as the problem solving algorithms) will prove VERY useful to PCT when we start exploring the control of higher level variables -- like programs and principles. I, for one, take my hat off to the ai people for their excellent accomplishments. They may not understand control or their place in it, but they do understand one level of the control model (the program level) better than any of us PCTers. We will be able to use their findings in our work eventually.

Regards Rick

Date: Mon Jun 08, 1992 11:23 am PST
Subject: Re: real and ideal

(sibun (920605.1600))

[Ray Allis 920605.1230]

we started out by my objecting to yr characterization of ``simulation.'' i think we will have to agree to disagree here. we seem to have gotten ourselves off on a tangent though, and i have some more fundamental problems with yr positions, which i don't think you've understood yet. i'll give it one more try below.

i still think you're contradicting yourself here, but it may just be a case of imprecise language.

you say:

> A (digital) computer running a program is affected by (physical) reality. The (physical instances of) symbols (as packets of electrons or whatever) are real and are affected by (physical) reality. But it's some electrons which are affected, not whatever it was that was symbolized by them. That relationship only exists in your mind.

great. i agree completely: the relationship between whatever the computer is doing and whatever you think it's doing depends on you.

you defined a simulation as something abstract and unaffected by reality. (you divide a simulation up into (object) symbols and relations, but that distinction is not relevant to my point.) a computer, down to its electrons, is squarely in reality. the simulation and the computer are only related via someone's mind.

assuming these definitions, it is ill-formed to say, as you did in a previous message:

> The computer operates on an arrangement of
> symbols, producing another arrangement of symbols.

the computer, being in reality, cannot be operating on the symbols, part of the simulation, which are outside reality. the computer pushes some bits around, and you interpret that as having some effect on the simulation. i think to preserve the distinctions you are making and to say what you want to say, it has to be something like:

> The computer operates on the bits in its memory and produces a different configuration of bits; we can interpret the change in this bit pattern as a change in the arrangement of symbols in the simulation.

--penni

They don't even have to symbolize anything (i.e., be symbols at all). The FORM's the thing.

We owe this sort of blind spot to the ancient Greeks.

i think the postulation of a distinction between form and meaning is more problematic than any ``confusions'' people have over the distinction. it focuses attention on the form rather than the reality that the form is deliberately divorced from.

Date: Mon Jun 08, 1992 11:23 am PST
Subject: Re: car event

(sibun 920607.1600)

From: David Goldstein
Date: 06/04/92

I was driving to work this morning. I was thinking about the discussion of modeling versus simulation, AI versus HPCT. I started to imagine putting a TV camera on the car which was part of a model designed to drive the car to work. I was wondering how the model would respond to a car which suddenly came close from the side direction. Then suddenly, just at that moment, reality intruded itself. The car on my left suddenly swerved in front of me. I quickly braked and avoided the collision. To the best of my recollection, this is the first time I was thinking about this sort of problem. It was kind of spooky. After the incident, the cars around me seemed to be doing maneuvers which seemed somewhat dangerous. This sense of danger continued for a few minutes until I exited the highway I was on. It was OK after that. Was this just a coincidence? Psychic phenomena? Was I in imagination mode and this played a part in the incident? What are your speculations?

i think you may have noticed the car was swerving toward you and that affected yr train of thought. you may have noticed some very subtle things about the car's behavior that never got as far as yr consciousness. i think a lot of what is considered psychic phenomena is the ability to notice things that others don't.

also, remember that you are *remembering* this, and you are stuck w/ yr reconstruction of the events. that is, the car may have perceptibly swerved toward you before you wondered about such a thing, even if that's not the order you believe (now) that the events happened in.

i was rearended a couple years ago. i remember the event as a long stretch of complete silence and stillness followed by an enormous jolt and noise. now, i don't think that during the few moments before the crash i was actually experiencing that silent stillness, even though my memory of it is very strong (i can easily evoke it now). presumably, i remember the stillness by contrast w/ the crash it preceded.

(btw, the car was totalled, but my passenger and i were essentially unhurt, due to impressive engineering and our wearing seatbelts!)

--penni

Date: Mon Jun 08, 1992 11:30 am PST
Subject: Recognition at last?

[Martin Taylor 920608 0930]

Headline to the lead editorial in today's Toronto Globe and Mail:

"Powers should follow purpose."

So what have you been doing for 40 years, Bill?

Martin

Date: Mon Jun 08, 1992 11:45 am PST
Subject: re: Standards

[From: Bruce Nevin (Mon 92048 12:49:17)]

(Martin Taylor 920607 1710) --

Welcome back, Martin.

The discussion of "standards" substituted that term for "principle" as in level 10. Standards meaning "norms" or "conventions" can be on any level. Modelling others to facilitate cooperative action with them involves perceptions on many levels.

The convergence of your discussion with the prior one is perhaps this: that people are aware of norms, conventions, and models of others mostly on the principle level, the level at which they attribute motivations and make moral judgments.

Bruce

Date: Mon Jun 08, 1992 11:58 am PST
Subject: A new control method

[From Rick Marken (920608 10:30)]

This is to Bill Chen.

Welcome to CSGNet. I have a couple of quesions about your post before I try to tackle the math (I don't think you'll find too many of us on CSGNet to be real familiar with many of the common tools of control engineering -- like Laplace trnasforms, etc. We tend to do most of our work with computer simulation and we are mostly interested in what variables living systems control -- but, of course, we are also interested in how they control them.)

I am currently working on a paper which describes a new method of aiding the human operator in a control task. So I want to understand your method. It looks to me you are using a type of predictive display system -- where the person is controlling a cursor relative to the a computed version of the target. The idea, I think, is to get the person to respond to this target (rather than the real target) so that their response will be "optimal" and make the actual difference between operator response and actual target be minimized).

Anyway, in order to see if I'm right I have to understand what you are referring to by some terminology -- CSG terminology is often somewhat different than that of engineering control theorists because we look at the control situation a bit differently.

So, you say:

>Where y() is the system output and u() is the control signal.
>At some time k, The time-series relation between outputs and
>control is

What do you consider the output signal and the control signal? Consider the task to be compensatory tracking carried out by a human subject (like keeping a cursor alighed with a fixed target). The cursor position, c(t), depends on handle movements, h(t), and disturbance, d(t), so c(t) = h(t)+d(t). Is h(t) the the output variable? I think it must be. But what is the "control signal"? is it d(t) (which determines the required value of h(t) that keeps the cursor on target) or is it the difference between target position (T) and cursor position, what is sometimes called the error input e(t) = T-c(t)?

>Consider the basic step-response control, the system outputs
>are always desired to approximate a fixed reference.

What is the fixed reference for in the compensatory tracking task -- for the difference T-c(t)? Who "desires" this fixed reference --the subject or the observer of the subject?

> On the
>steady state, the system outputs are equal to the reference
>and the control signals are usually kept at a fixed value.

If system outputs are h(t) then the only reference they should be equal to is d(t). So this would mean that the reference is to have system outputs equal to control signals?

It is important to know what you think is a "reference" and who sets it, because you say:

> An interesting skill : for varying references, if the future
>   references are known at every sample, the error function can be
>   modified to

```
>           /  y(k+i-n')  \   /  r(k+i-n') \
>           | y(k+i-n'+1) |   | r(k+i-n'+1)|
>           |      .      |   |      .     |
>           |      .      |   |      .     |           1+a1+...+an
>       E = |    y(k+i)   | - |   r(k+i)   |   us(j) = ------------- * r(j)
>           |  u(k+i-m')  |   | us(k+i-m') |           b0+b1+...+bm
>           | u(k+i-m'+1) |   |us(k+i-m'+1)|
>           |      .      |   |      .     |
>           |      .      |   |      .     |
>           \    u(k+i)   /   \   us(k+i)  /
>
>    This skill results in a null-phase-delay response.
```

I presume that the r() are the reference values at different times. How do you know what these values are? Are they really the values that someone thinks that the subject SHOULD maintain as their reference for the difference between output and control signal?

I want to understand this because one of the main virtues of my own approach to aiding a controller is that it works even though the controller him/her self changes his or her OWN reference for the difference between output and control signal in an unpredictable manner.

Best regards Rick

Date: Mon Jun 08, 1992 12:31 pm PST
Subject: Re: Plasticity; history determined; miracles

[Martin Taylor 920608 12:15]
(Bill Powers 920523.0800 -- long time ago, but that's where I've got to)

(Bill on Greg Williams 920522-2)
>I'm not deliberately misreading you, but perhaps I'm misreading you. To me,
>"history-determined" doesn't seem to mean what it means to you. I don't
>think that "history," which is simply a record of what we remember of the
>past, has any influence on anything (although our memories of it, which
>exist in present time, may have some influence on what we do next). Even
>the FACT that something happened in the past has no physical effects NOW.
>All physical effects, I assume, occur in present time, as one variable
>influences another coexisting variable. The past can't affect the present
>any more than the future can. This doesn't keep us from tracing a course of
>events through time, but all that happens in imagination.

In a trivial sense, this must be true. Interacting events are on the same light cone. But I don't think it helps us to understand what is going on. For example, if I am right about the relation of sense-data to categorical perception, history matters very much. Categorical perception depends on the existence of (at least) fold catastrophes, and which sheet of the fold defines the NOW percept depends on history as well as on present sense data.

Control depends on NOW data, but the bandwidth of the system determines the duration of NOW. If the bandwidth of the control system is measured in cycles per day, then NOW extends over substantial fractions of a day for that control system (Nyquist). Again, history matters, because what a low-level (high-bandwidth) ECS is doing depends on the perception of a higher-level one that might informationally be quite old.

(I do realize that bandwidth concepts become quite hairy when we are dealing with non-linear systems, especially when their control surfaces contain catastrophes, but I think the ideas remain conceptually valid except in discrete cases when things change fast).

> A calculation of voltage
>from resistance times current, for example, doesn't depend on how a
>particular resistor happened to be manufactured, or on what is supplying
>the current, or on what the current was prior to the measurement.

You are talking about a linear system whose state EXPLICITLY does not depend on history. If you talk about a magnetic system with hysteresis, you might word it a bit differently.

On a different issue, you said in one posting sometime earlier that in a chaotic system you can work from the present to the past, but not to the future, because the dependence on initial conditions implies a divergence into the future. You assume that this entails a convergence into the past, but in many chaotic systems this is not the case. The divergence goes both ways--many pasts can lead to almost the same present state, just as many almost identical present states lead to different futures. We cannot deduce history.

Martin

Date: Mon Jun 08, 1992 1:30 pm PST
Subject: Control-sharing; reinforcement

[From Bill Powers (920808.1400)]

Martin Taylor (920606.0815) --

Your design of a "control-sharing" system is truly ingenious. I like "insistence" for "gain." In other circumstances, "importance" works well, too.

In light of Rick Marken's "conflict stabilizer" system, it may be a good idea to leave certain low-level controlled variables under the automatic system's control all the time, to get around the pilot's transport lag.

Any chance of your coming to our meeting July 29-Aug. 2?

Oded Maler (920605) --

>Suppose an agent is in a closed-loop relation with the "world". No
>matter how it is organized inside, what it perceives of the world
>results in the "actions" it performs (e.g., muscle contraction) which
>again, cause certain things in the world, affect the agent's
>perceptions and so on. "Reinforcement" in this context speaks about how

>this perception-->action map changes with the history of interaction.

I don't object to your usage of "reinforcement," but I'd rather drop the term because of its history. It's been used uniformly as if reinforcement were an effect of some external thing or situation that shapes the behavior of an organism. Isn't it just as informative to speak of controlled variables and reorganization?

Your way of putting matters gets uncomfortably (for me) close to sequential analysis of a closed loop, as if a whole perception takes place, as an event, and then an action-event occurs that affects the next perception-event, and so on, in alternation. The sequential nature of language makes this difficult, but I think we should always try to emphasize that actions and perceptions vary concurrently, changing at the same time and not taking turns. Without this concept it's difficult to see that the action is controlling the perception.

Also, the sequential kind of description makes it seem that a perception causes an action, which is like the old SR idea. When you leave out the reference signal, you imply that a given perception will always be associated with the SAME action, where in fact it can be associated with completely opposite actions (if the reference signal rises or falls). It DOES matter what happens inside the organism!

My main point about reinforcements, as events, things, or situations, is that there is always an implicit reference-level for the event, thing, or situation. The organism acts on the environment to make the experienced events, things, or situation match the internal reference states for those things, to the extent it can. If you get cause and effect backward, it seems that the increase or decrease in the reinforcer is causing the behaviors, where in fact it is the behaviors that are bringing the reinforcer nearer to its reference level, via the external feedback link or "schedule." I think the potential for confusion in this term is just too great.

Best, Bill P.

Date: Mon Jun 08, 1992 1:37 pm PST

Bill Chen (920608) --

Hello, and welcome to CSGnet.

RE: New control method.

I think that you are defining "control" quite differently from the way I do. It seems that you are computing an output y(t) that is a function of an input or "control" u(t) where the function involves lagged values of u(t). The purpose of the computation appears to be to "forecast" future values of y, given present and past values of u and past values of y. The forecasting appears to be used in making future values of output y match some reference values that are known in advance.

It may be that you have an elegant alternative to the inverse-kinematics approach -- I am not familiar with z-transform methods, but I will take your word that the derivations are correct.

Unfortunately, even if your method works it will not explain the behavior of organisms, the main concern of CSGnet (neither will any inverse-kinematics model). Organisms do not produce specific outputs, but specific outcomes. The arrangement is this:

```
      ORGANISM ---> OUTPUTS ---> OUTCOMES OR BEHAVIORS.
                                          ^
                                          |
             INDEPENDENT DISTURBANCES ----
```

The actual outputs of an organism are (mainly) its muscle forces. The outcomes or behaviors are the effects of these muscle forces on the environment PLUS the effects of independent variables -- disturbances -- that are part of every natural environment. When we observe regular or repeatable behaviors, we are not observing muscle forces, but the vector sum of muscle forces and independent disturbances. So there is no regular or unique connection between the organism and the effects we call behaviors.

The only model that can explain how the outcomes can remain regular despite the disturbances is the classical closed-loop control system model. The organism must be SENSING the outcomes or behaviors, comparing the sensed outcome with some reference state, and adjusting its outputs on the basis of the difference. The "reference" values in question do not define desired states of the output, but desired states of the sensed outcome.

I should point out that in general the disturbance variations are unpredictable; in most cases the causes of disturbances are not sensed by the organism, or if sensed are not sensed with anything close to the accuracy needed to explain the stability of the outcome. The classical control model shows how the outcome can be controlled even when the organism has no information at all about the causes of disturbances, and is incapable of predicting their future behavior. I do not believe that your model could work under such conditions. When people on this net speak of "control", they mean control of outcomes, not outputs.

Best, Bill P.

Date: Mon Jun 08, 1992 1:41 pm PST
Subject: A problem of convergence

[From Bill Powers (920608.1400)]

David Chapman (920606) and Phil Agre -- CC: CSGnet

I've read some papers that Agre sent me (much thanks!). I'm left wondering what to do. If I come over and try to play your game on your field, I'll be too ignorant to say anything interesting to you. If we stand on our respective fields and throw fastballs at each other, it won't be much of a game. Although it would be a dirty shame, it's possible that our worlds simply don't intersect over a large enough area to support a dialogue. I hope that this isn't true.

I think that HPCT redefines some fundamental problems in modeling behavior. It doesn't substitute for analysis of symbol-handling systems, and it certainly has some family resemblances to the guts of Pengi etc. But I just know too little of your interests to serve as an interpreter in any way that would be useful to you. You're hell-bent-for-leather in a certain direction, and if I try to get in the way I'll just get trampled.

At the risk of sounding standoffish or condescending, I think there's only one practical way to find convergences between your approach and mine, and that is for you (two) to learn HPCT. If you could take the time to pretend to be students again for a while, and learn the basic principles behind HPCT, you would be by far the best judges of its meaning in your own work. I know how that sounds and I apologize for the implied hubris. But I just can't think of any other way to go about this that won't lead to endless misunderstandings and arguments about words. In Agre's articles I can see opportunities for hundreds of frustrating interchanges. But I would greatly value your participation on CSGnet if we could bypass all that stuff.

As to your suggestions about possible publication of the Gatherings simulation, thanks, David, but it's moot. A paper on this model by McPhail, Powers, and Tucker will appear as the lead article in the next issue of Computational Sociology (if I have the title right).

Best, Bill P.

Date: Mon Jun 08, 1992 1:54 pm PST
Subject: A problem of convergence

Well, in fact my interests have gone off in entirely different directions since the Pengi work, and really I don't think about that stuff any more. (I think the same is true of Phil.) There's still plenty of people in AI interested in these sorts of issues, but we aren't them. Sorry about that...

Date: Mon Jun 08, 1992 2:11 pm PST
Subject: History determination

[From Bill Powers (920608.1600)]

Martin Taylor (920608.12.15) --

Making your way toward present time, I see.

>Control depends on NOW data, but the bandwidth of the system determines
>the duration of NOW.

A nice point: the "specious present" seems, even subjectively, to get longer as you consider higher-order variables.

My point concerned modeling more than explanation (understanding what is going on). Think what's involved when you just try to say that y[t] =

f(x[t-1]) in a computer simulation. You can't actually make y depend on the value of x from a previous iteration -- unless you SAVE that value of x for use in the next iteration. This is even more obvious in analog computing, where you can't save any past values of any variables. Everything that makes the system work has to be assembled in the NOW if any interactions are to take place. I think this is true of the operation of any real system: literally, the past is gone unless it is specifically preserved as a present-time effect or memory.

When it comes to explanation, on the other hand, it's a different matter. In an explanation, we try to explain the time-course of processes, tracing them from one moment to the next (or backward to a previous moment). Now you can look at the charge on a capacitor and say how it got to be that way -- through integration of current. As far as any PRESENT effects of that charge are concerned, however, it doesn't matter how it got that way. Only the present state matters in determining what will happen next.

Even fold catastrophes present the same distinction between modeling and explanation. To explain the current post-catastrophe state, you consider not only what did happen but what might have happened if the state had been reached by a different path. But at each successive NOW, it's only the current state that matters in the interactions NOW taking place. At each moment, the variables are on a particular path: when a bifurcation comes along, its knife-edge exists NOW, and the result will be only one next state. If you include all relevant derivatives in the NOW, history literally makes no difference. All real interactions take place only in the present. Even magnetic hysteresis.

I did manage to realize that chaos implies an uncertain past as well as an uncertain future -- you probably haven't got to that one yet. I was originally thinking mainly of the problem of hypersensitivity to initial conditions in predictions involving integrations. I didn't mention, by the way, dissipative systems, in which the future state is quite predictable in general terms: the marble will eventually, by some path, come to rest at the bottom of the bowl. But these are special cases applying mostly to the inanimate world. A dissipative system with a constant renewal of the energy supply is a different beast.

Best, Bill P.

Date: Mon Jun 08, 1992 2:13 pm PST
Subject: Control-sharing

[From Rick Marken (920608.1500)]

Bill Powers (920808.1400) says:

>Martin Taylor (920606.0815) --

>Your design of a "control-sharing" system is truly ingenious.

I agree. I appreciate being mentioned as the inspiration for it. But your proposal (Martin) is far more ambitions than my "conflict stabilizer" system for improving human control. I think your system sounds like a real intelligent step towards the appropriate integration of people and machines. The theory sounds ok; now all (?) you have to do is develop the technology that will do little things like reliably determine when the operator's gain or reference is changing. Sounds like a research project that could keep some graduate students busy for a few years.

Welcome back. Rick

Date: Tue Jun 09, 1992 8:35 am PST
Subject: Re: habits

[Martin Taylor 920609 12:30]
(Andy Papanicolaou 920603 15:31)

I hate to be repetitive or (as Rick Marken delights to be) predictable about this, but on the learning of perceptual skills and motor habits...

>We can go on and on with difficulties, but at some point, we must
>begin drawing the outlines of a CT model that simulates the process
>of this sort of habit formation. If you have any suggestions, Tom
>and I would surely appreciate them.

J.G. Taylor. The Behavioral Basis of Perception. Yale UP 1962.

JGT predicts exactly the phenomena you describe about learning foreign phonemes, as well as much else that seems a necessary framework for how the control hierarchy comes about, though he does not describe the control hierarchy itself. A difficult book to read, but worth the trouble.

Martin

Date: Tue Jun 09, 1992 9:00 am PST
Subject: Re: Models & simulations; habits

[Martin Taylor 920609 12:35]
(Bill Powers 920603.1700 responding to Andy Papanicolaou)

>Didn't Columbus think that if he sailed West, he would reach Cathay?

He would have, except he got the diameter of the Earth wrong and would have starved first if something unexpected hadn't got in the way.

>You can't control the sounds of "l" or "r" reliably if you can't perceive
>the difference between them -- auditorily and kinesthetically/tactily. If
>you have the wrong feel going with the right sound, you'll feel the wrong
>articulation and hear its result as right.

Do I detect linear cause-effect thinking here? JGT's results that you can learn to perceive only what you actively control would argue that the learning is itself a feedback loop. You learn to perceive by doing, and you learn to do by perceiving. It doesn't work one way OR the other. You are buiding an ECS (I would guess) for the percept, and only by controlling it incrementally better do you learn to perceive it.

And if I can add a personal belief to that: you learn the new percept better, the more variable the means you use to achieve it. By varying the actions, the appropriate percept is better isolated from the percepts associated with the lower-level control systems involved in its control. From a quite different standpoint, I came to the same conclusions in respect to second language learning in my "Speculations on bilingualism and the cognitive network"
(Ontario Institute for Studies in Education, Working Papers on Bilingualism, 1974).

Martin

Date: Tue Jun 09, 1992 10:18 am PST
Subject: blindmen

[From Rick Marken (920609.1100)]

Well, the revised copy of the "Blindmen.." paper is ready. I sent a version to Psychological Science a couple weeks ago. Estes liked it but was reluctant to publish it as a seperate theoretical note (I basically agree with the rationale he gave; Estes is definitely a decent person). He encouraged me to submit is as a reply to another article; I might do that but I've decided to sent it first to Psychological Review to be published as a theoretical note (I think they publish those?). If anyone wants to see the revised version I would be happy to sent it by personal e-mail; then you can see all the advice I didn't take (actually, I used most of the suggestions I got).

Hasta Luego Rick

Date: Tue Jun 09, 1992 10:46 am PST
Subject: Re: AI, PCT, and HPCT

[Martin Taylor 920609 14:20]
(Bill Powers 920604.0800 responding to Penni Sibun 920603.1600)

(Sorry for the flood. I'm trying to clear off my mail backlog today, most of it from CSG-L).

I'd like to reinforce Bill's comment "if you send the same driving signals to the muscles twice in a row, you'll be lucky to see any resemblance at all between the behavioral outcomes on the two occasions." Recently I have been using a simple example in communication to illustrate this. It's the inverse of the usual illustration of control, which says you get the same result by variable means.

The example is that person P says to person Q "Can you close the door?" In each example, the intonation is the same, but the meaning is very different.

(1) Normal: P is sitting down, Q is by the door. P wishes to perceive that the door be closed, and P believes Q to be able and willing to do it. This

situation is ordinarily called an "indirect request."

(2) P is a physician, Q has been suffering from some muscular disability. P wants to ascertain the extent of Q's disability. Q's appropriate response is "Yes" or "No", and P is not controlling for a perception of the door being closed.

(3) P has tied Q up so Q cannot reach the door. P is controlling for the perception of Q feeling humiliated.

(4) P is a high-level boss, Q has been asked to come to P's office. P may not be controlling for the perception that Q should be terrified, but it is a likely consequence of the question. P is controlling for both the perception
of power over Q and for the door to be closed.

(5) Q has been boasting about athletic prowess. P uses the question as an insult, to deny that prowess. P is not controlling for the perception of the door being closed, but is controlling for perception of signs of annoyance from Q.

I think that the "meanings" of the exact same acoustic waveform are sufficiently
different as to admit of almost no common core. Different perceptions are being controlled by identical means. Different information is being transmitted from P to Q by the same "code" in the different situations.

>If you ask how it can be that variable means produce consistent and often
>closely controlled and disturbance-resistant outcomes, you end up with
>control theory. That's the ONLY explanation anyone knows of that works. PCT
>isn't optional; it isn't just an alternative view. It's the inevitable
>result of admitting that outcomes are in fact under control (meeting a
>formal definition of control), and seeing, eventually, that the only way
>to explain this fact is that the organism is controlling its own
>perceptions of those outcomes.

Amen. It's the way this truth is developed that distinguishes possible theories. Without this truth, no behavioural or psychological theory can be viable.

Martin

Date: Tue Jun 09, 1992 11:29 am PST
Subject: Re: History determination

[Martin Taylor 920609 15:00]
(Bill Powers 920608.1600)

>Making your way toward present time, I see.

Almost here...

> I didn't mention, by the

>way, dissipative systems, in which the future state is quite predictable in
>general terms: the marble will eventually, by some path, come to rest at
>the bottom of the bowl. But these are special cases applying mostly to the
>inanimate world. A dissipative system with a constant renewal of the energy
>supply is a different beast.

Yep. Might even be a control hierarchy. At least all control hierarchies have to be dissipative systems with an energy source and sink. Don't forget the garbage/shit. It's very important to the system. You can't have life without it.

A general point. There really isn't a sharp distinction between, on the one hand environments that are absolutely stable and in which prediction of the consequences of action are predictable, and on the other environments that are totally disturbed and in which no prediction is possible. For the most part, we could anticipate that our actions have consequences that would fall in
some range even if we did not perceive them. Predictive modelling as an aid to control--planning--is not stupid in such an environment. I realize that several of the strong comments antagonistic to prediction and planning that have been made by Bill, Rick, (and others?) over the last few weeks have been aimed at people who do not appreciate the centrality of the NOW control of perception, but the way these comments are sometimes phrased seems to me unfortunate.

As a low level example, linear predictive coding is a good way to reduce the information required to identify the NOW state of a speech sample, by predicting what it should be expected to be, given the last several samples. The prediction is hardly ever exact, but it's much better than saying "that there's speech, that be; tell me what this sample be." Most of the world is somewhat predictable, most of the environments in which we want to achieve certain perceptions are moderately stable. Habits, plans, and predictions can ease and make more precise the normal work of control, at the cost of making control harder when surprising disturbances do occur.

Martin

Date: Tue Jun 09, 1992 12:28 pm PST
Subject: VOR Disorder

[from Gary Cziko 920609.1500]

To Wayne Hershberger and any other brain-eye experts:

A relative of mine has been having dizziness/vertigo/balance problems. The patient's doctor originally thought this was caused by excess pressure in the inner ear which disrupts the vestibular system (in acute stages called Meniere's disease), but now he thinks it is more central and that the problem is due to a malfunctioning of the vestibulo-ocular reflex (VOR). The patient reports a world which is never completely visually stable, especially after quick head movements and after periods of activity which result in head movements causing everything to "swim."

I am hoping that someone on the net (like Wayne Hershberger) could give me a brief explanation of how the VOR works (particularly from a PCT perspective, although I understand that at some basic level this is supposed to be an open-loop system) and point me to some references that would give me some more basic info. The doctor has said that there is not much he can do about it, but recommended that the patient not try to avoid conditions that cause the problem since he feels that this is the only way (by "stressing the system") that it has a chance of fixing itself. Perhaps PCT can tentatively offer some novel ways of approaching this problem. Meanwhile, I'll be checking through the medical literature to see what I can find.--Gary

Date: Tue Jun 09, 1992 6:48 pm PST
Subject: Perceiving vs. Controlling

[from Gary Cziko 920609.2100]

Martin Taylor (920609 12:35) said:

>JGT's results that you can
>learn to perceive only what you actively control would argue that the
>learning is itself a feedback loop. You learn to perceive by doing, and
>you learn to do by perceiving. It doesn't work one way OR the other. You
>are buiding an ECS (I would guess) for the percept, and only by controlling
>it incrementally better do you learn to perceive it.

But can't one learn to perceive all kinds of things one can't control? A speaker of North American English can understand without much difficulty all kinds of worldwide dialects and accents of English but couldn't come close to sounding the same way. I can tell (usually) an Australian from an Englishman, but there is no way I could imitate their "accents" so that someone else could tell which one I was trying to produce.

So while I can understand Bill Powers's statement that can't control if one can't perceive, but I don't understand the claim for the converse. What am I misunderstanding about JGT and this claim?

--Gary

Date: Tue Jun 09, 1992 7:31 pm PST
Subject: JGT

[From Rick Marken (920609.2030)]

Well I'm glad someone noticed this again:

Gary Cziko 920609.2100 says:

>Martin Taylor (920609 12:35) said:

>>JGT's results that you can
>>learn to perceive only what you actively control would argue that the

>But can't one learn to perceive all kinds of things one can't control?

It seems so to me too. In fact, if JGT were right, I'd be functionally blind. I can't control my kids, my wife, my mother. How did JGT come up with this stuff?

Maybe if he'd left that part out of his theory he might be a household word today. Even Freud didn't say anything THAT silly.

Confused in California Richard S. Marken

Date: Wed Jun 10, 1992 5:47 am PST
Subject: MMT on JGT

(Rick Marken (920609.2030)) --
(Gary Cziko 920609.2100) --
(Martin Taylor (920609 12:35)) --

>>JGT's results that you can
>>learn to perceive only what you actively control would argue that the

>But can't one learn to perceive all kinds of things one can't control?

My take on this: In translating JGT's work to CT terms, MMT has used the word "control" where it doesn't belong.

Here's a stumbling try to make the relevant distinction: JGT showed that you can learn to perceive only what you can interact with, where "interact" entails a feedback relationship through the environment, but where "feedback relationship" does not always entail control.

Maybe someone else can get a cleaner cut at that.

Bruce bn@bbn.com

Date: Wed Jun 10, 1992 7:05 am PST
Subject: Habits and speech

[From Bill Powers (920610.0800)]

Chuck Tucker: $90 should cover it, including membership, banquet for both, no other meals for you or guest. You can purchase meals with the rest of us, as desired (Breakfast $4.50, Lunch $5.50, Dinner $6.25).

Andy Papanicolaou & Tom Bourbon (920606) --

RE: habits and speech

I guess I misunderstood the thrust of your questions about habit formation. But I still think the problem lies at least PARTLY in focusing too much on outputs. I don't mean to dismiss the fact that output learning (learning of lower-level control systems) must take place, only to point out that many output processes are as they are only because the physical extra-neural world is as it is.

>So, we believe that it is generally true (and, hopefully, the linguists
>will concur or persuade us otherwise) that every time a /ba/ or a /ga/ >or
a /tu/ is heard the corresponding patterns of articulatory gestures
>contain a set of invariant features.

OK, this is true, even if it is also true that there are many sounds that can be made by producing variable articulator patterns. Consider, however, the role of the non-neural environment in the way I open the front door of my house when entering from outside. It is my habit to grasp the knob with my left hand, turn the knob clockwise, lean slightly toward the door while turning the knob and pushing, and extend my arm so that the door swings open, inward. There are slight variations in this habit, but most of the time that is how it's done.

The question is, how does this habit come into being, and why this one instead of some other? We're all agreed, I think, that to answer "Because the opening of the door reinforces the movements that lead to its opening" would leave us no closer to an explanation. What we would like to know is how I come to do just those acts that will open the door, and why I do them.

The most general answer is "Because I want to get inside the house." The reference signal that motivates the opening of the door is a picture of myself stepping through an opening that does not initially exist. If I began this process in total ignorance, I might try pushing the door out of the way, pulling it out of the way, sliding it out of the way, waving my hands and crying out "Open, Sesame!", or walking back and forth in a figure-eight pattern. None of these actions, however, would be effective -- not because they are foolish or superstitious, but simply because they don't have the effects on the door that suffice to open it. When one has no knowledge of the properties of the door, there is nothing to say that one act would work any better than another. Trying one action that one knows how to produce is no more foolish than trying another.

If I'm solving this problem by reorganization -- by running randomly through different uses of skills I already possess -- I will simply keep trying until I hit on the combination that makes the actual perception (a closed door) change to match the desired one (an open door).

If I'm solving this problem by reason and insight (in the way a Djinn would who is used to doors that open with "Open Sesame"), I will see that there is a knob, and devices holding the door on its left side that look as if they could pivot and would prevent sliding. By examining the door frame, I will see that it is prevented by a strip of wood from opening outward, and deduce that it must be pushed inward. I will realize that the knob, being rotationally symmetrical, is probably intended for grasping and turning. I will see that I should push the knob with my left hand so my arm will be out of the way if or when I move through the door. I will then end up doing exactly the same actions I would have done at the end of a random search for an effective means.

The fact is that the door is so constructed that only a combination of knob-turning and pushing will make it swing open and create the wanted perception of an opening sufficiently large to walk through. That fact is quite independent of the organization of my nervous system. Doing this with the left hand will leave me free to move through the door while keeping control of it. It is the physical door and its properties that determines what lower-level variables I MUST control in order to gain control of the higher-level variable that is my reason for wanting to open the door. While there are slight variations in the way different people might go through this process, everyone who opens that door will end up doing it in essentially the same way -- not because there is any propensity of people to develop similar habits of acting, but because there is a limited number of actions that will open the door. Given the purpose of perceiving an open door, and the initial condition that the door is closed and latched, essentially everyone will perform the same acts in opening it. This tells us nothing about the nervous system, but much about the door.

Now transfer this parable to the pronunciation of certain sounds. If, in fact, there is only one manipulation of the articulators that will produce a given sound, /ba/, and given that a person wishes to experience the sound of /ba/ (perhaps simply in imitation of a remembered sound), then however one succeeds in doing this, by random trial and error or through visual cues or by clever reasoning, the final result will be to close the lips, start the sound, and release the lips. This is not because of any propensity to develop that habit, but because that is the set of physical processes that produces the wanted sound. If different people end up with the same articulator habits in making the sound /ba/, this is not because their nervous systems have similar inclinations to develop that habit, but because the physical construction of different people is, in the relevant regards, very similar. That is where the invariant features of "articulatory gestures" come from.

Suppose that there is more than one way to open the door: suppose it can swing open either inward or outward, also also slide sideways into its frame, after the knob is turned either clockwise or counterclockwise. Starting in ignorance, I might try just pushing and pulling, or walking in figure eights. Eventually, I will hit on a combination of acts that creates the perception of an open door that I want. This combination then becomes a means of opening the door. I turn the knob counterclockwise and slide the door open. Why, then, do I use this act again the next time I open this door? Why does this act become a "habit" even though other acts would work just as well?

The answer is simple: if the higher order goal is simply to get through the door, and the output of this higher order control system has become connected to reference inputs of lower systems so as to create the desired perception, there will be no error to motivate another random search (or another period of clever reasoning). The alternate solutions will go undiscovered because there is no need for them. The "habit" of accomplishing the goal by one specific means remains by default. Only if this means entails awkwardness or pain or some other error-creating consequence would there be any motivation to extend the search.

Now consider the Japanese who has trouble producing an 'l' sound, but is quite capable of producing an 'r'. The goal, in a Japanese environment, would initially have been to hear the same 'r' sound heard in others' speech. So the reference signal would have been this remembered sound. The Japanese baby, by experimentation, finds a way to manipulate the articulators so they create the desired sound. After much learning, the feeling of making the sound is added to the auditory experience, and remembered, so the effective reference signal includes kinesthetic as well as auditory components of lower order.

Now the sound 'l' is heard, and for some reason the (now-older) Japanese wishes to reproduce it. Initially, only the 'r'-input-function responds to the heard sound of 'l'. So the Japanese understands that the goal is to produce the 'r' sound, and does so. "No, no," the Western teacher says, "You must say it this way: 'l'." And the Japanese, hearing 'r', says 'r'. This could go on forever: it seems that the Japanese pupil simply has too strong a habit of making the articulations that result in the sound of 'r'. In fact, there is no "habit." The Japanese is simply repeating back the sound that is heard.

The Western teacher, who has different perceptual functions that respond to 'l' and 'r', concludes that the Japanese is lacking the ability to SAY the sound 'l'. In fact, what is lacking is the ability to perceive the sound 'l' as something different from 'r': both are heard as 'r', although the 'l' might not generate quite an optimal sense of 'r'. The puzzled Japanese hears himself or herself reproducing the same sound that the teacher is emitting, and doesn't understand what the problem is. From the Japanese viewpoint, neither the Japanese nor the English language contains the sound 'l'.

If, by some means, the Japanese could be persuaded to search for perceptual differences between the two sounds, and ultimately to be able to indicate which sound is being heard, the problem would still remain of creating those sounds -- just as a person who perceives that a door can swing open as well as slide open must still learn to use different actions to create this new perceptual consequence. Now kinesthetic learning must take place. The new perceptual function (and comparator) must produce an error signal, and that error signal must be connected to produce new reference signals for the way the mouth and tongue feel while uttering a sound. If the Japanese understood that this is now the problem, he or she could begin experimenting until the sound was successfully produced. Then, of course, a period of practice would be needed so that the initial sound-only perception became a sound-plus-feel perception, the way the 'r'-sound controller has developed.

It occurs to me (thanks to Martin Taylor and his unrelated namesake J. G. Taylor) that the perceptual learning might be facilitated if the person were given control of the sound -- not through the articulators, but through some simple means the person can already use, such as a control handle. On hearing the teacher say "road" and "load" over and over, the person would hear an artificially-produced word "road" being repeated over and over, with the initial consonant being adjustable smoothly between the sound of 'r' and the sound of 'l'. The task would be to move the handle until the artificially-generated word matched the sound of the word as pronounced by the teacher. Once control of the auditory 'r'-'l' difference is attained by this means (assuming it ever is), perhaps the task of finding a different means -- using normal methods of articulation -- would then be mastered more easily. Martin Taylor has pointed out repeatedly that perceptual learning works best in a control context -- perhaps ANY control context is better than none.

Marvin Brown says that once a wrong pattern of articulation is learned, it is too late to learn native speech. Perhaps the stage of using an alternative means of producing the right sounds might overcome this problem.

Marvin received your message, Gary, and will probably reply by snail mail -- I have to conclude now to have a final conversation with him and wave goodbye as he heads back to Salt Lake City.

Best, Bill P.

Date: Wed Jun 10, 1992 10:35 am PST
Subject: JGT

[From Rick Marken (920610.11:00)]

Bill Powers (920610) has helped me understand what Martin and JGT might have meant by the claim that 'you can learn to perceive only what you can actively control'.

Bill says:

>It occurs to me (thanks to Martin Taylor and his unrelated namesake J. G.
>Taylor) that the perceptual learning might be facilitated if the person
>were given control of the sound

This is a very interresting suggestion. I believe that one does develop new perceptual functions in order to develop better control. I think that is a large part of what skill learning is about -- learning what to perceive (and control) in order to control a higher order variable. I am sure that a skilled pianist perceives relationships between the black and white keys (when they are playing in a particular key) that I currently don't perceive. Having the demands of controlling some variable might push the (re)organization of new perceptual functions, so that now you can perceive a variable that had not previously been perceived. Maybe all perceptual functions ARE developed through efforts to control -- so JGT is right. And the reason we can perceive so many "things" that we can't control is because we are seeing the world in terms of the perceptual functions that we have developed through out efforts to control. For example, as a child we might have developed perceptual functions that let us perceive "flatness" and "distance" and "height" and "greyness" as a result of our efforts to control our relationship to surfaces, objects, and whatever. Once we have these percetual functions available we then perceive in terms of these functions -- even when we are not controlling. So we can perceive the flatness of the walls or the height of a building even though we cannot (or do not) usually control these perceptions of the building.

If this is what JGT is saying, then I guess I agree; it may be true that we only develop perceptual functions though efforts to control; a process that requires the ability to perceive in terms of particular functions of the "outside world". Once these functions have developed, though, we can perceive in terms of them, whether we are controlling or not. A superficial reading of the claim that you can only perceive what you can actively control leads to the impression (which Gary and I had) that JGT claimed that you can only perceive what you are actively controlling. This, of course, is not true and, apparently, not what JGT meant.

No longer confused (about JGT) in California

Rick

Date: Wed Jun 10, 1992 10:35 am PST
Subject: Re: Perceiving vs. Controlling

[Martin Taylor 920610 13:30]
(Gary Cziko 920609.2100, Rick Marken 920609.2030, Bruce Nevin 920610 0908)

Gary:

>But can't one learn to perceive all kinds of things one can't control? A
>speaker of North American English can understand without much difficulty
>all kinds of worldwide dialects and accents of English but couldn't come
>close to sounding the same way. I can tell (usually) an Australian from an
>Englishman, but there is no way I could imitate their "accents" so that
>someone else could tell which one I was trying to produce.
>
>So whil I can understand Bill Powers's statement that can't control if one
>can't perceive, but I don't understand the claim for the converse. What am
>I misunderstanding about JGT and this claim?--Gary

There's no problem here, except the oft repeated one of mistaking a reference value for the existence of an ECS that uses that reference. The percepts that allow you to distinguish accents or dialects are the same (but with different values) as those that allow you to produce your own accent and dialect. No new kinds of ECS need be developed ("kind" means having a specific
perceptual function transforming specific inputs into a value that can be compared with a reference value).
=============
Rick:

>>But can't one learn to perceive all kinds of things one can't control?

>It seems so to me too. In fact, if JGT were right, I'd be functionally
>blind. I can't control my kids, my wife, my mother. How did JGT
>come up with this stuff?

>Maybe if he'd left that part out of his theory he might be a
>household word today. Even Freud didn't say anything THAT silly.

>Confused in California

But you can control for perceiving the presence of your wife (as opposed to a generic woman). Again, we have the distinction between reference values and ECS types. And I don't remember JGT EVER saying anything really silly except on purpose (which he did quite often). You ARE functionally blind

for the kinds of things you haven't learned to control for. Can you, with your eyes closed, identify the material, the location and the size of a rectangular object 15 cm high placed on a table at arm's length in from of you? I doubt it, and the reason is that the perception involves control for something of a kind new to you. But JGT showed you can learn this perceptual ability, if you learn control of it. Having someone tell you what is there is no good. Reaching out and feeling while you are hearing the echoes from the object is one way to learn (perhaps the only way). When you do that, you are learning by developing a control system that would permit you to select for perception of particular objects.
================
Bruce:

>My take on this: In translating JGT's work to CT terms, MMT has used the
>word "control" where it doesn't belong.

>Here's a stumbling try to make the relevant distinction: JGT showed
>that you can learn to perceive only what you can interact with, where
>"interact" entails a feedback relationship through the environment, but
>where "feedback relationship" does not always entail control.

There's a subtle thought. I think Bruce is right, in a way. Initially, when you are learning to perceive something of a new kind, you cannot control. What you are doing is learning to control. I think you are building a new ECS, but equally it might be changing the function of an existing ECS that had a different but related perceptual input function. I have to acknowledge once again that JGT did not conceive things in terms of the same kind of control hierarchy that we take for granted in this group. Sometimes I put my own interpretation in PCT terms on what he actually said, without realizing that I am doing it.
=====================

General comment:

Of what use is a percept that cannot be controlled?

The resources devoted to creating the percept must at best be excess baggage when it comes to the survival or effectiveness of the perceiver, and at worst the uncontrollable percept may be distracting and dangerous. At the same time,
it is equally clear that almost all perceptions are, at any moment, not being controlled. This is the reason for alerting functions (such as the movement detectors in the visual periphery), and is at the base of concepts such as "situation awareness."

Gary and Rick mistake the perception of variations within a type for the perception of new types. An ECS is concerned with the variation within a type (be it sensation or principle), but it cannot handle sensory abstractions that are of a different type. I take "learning to perceive" as "developing an ECS that can work with variations on a type." That, within H-PCT, implies that the ECS can be provided with reference values. Learning what can be done with the new percept is something else again. So, as Bill says, learning to control requires that there be a percept to be controlled; and I say, learning to perceive requires learning to control that perception.

Maybe this doesn't clear up the question. But it may help.

Martin

Date: Wed Jun 10, 1992 12:07 pm PST
Subject: Re: Perceiving vs. Controlling

[From Rick Marken (920610 12:30)]

Martin Taylor (920610 13:30) says:

> The percepts
>that allow you to distinguish accents or dialects are the same (but with
>different values) as those that allow you to produce your own accent and
>dialect.

This says it very nicely. As you can see from my earlier post today, I caught onto this after embarassing myself by suggesting that JGT might be saying something silly. In fact, it seems that he is saying something quite deep; philosophically, I suppose, JGT is an empiricist -- he's saying that we LEARN to perceive (perceptual functions are not built-in at birth). I think Bill P. made a similar suggestion in BCP -- when he suggested that a perception like "straightness" could be learned by creating a perceptual function that produces a signal that is invariant with respect to certain outputs (such as translational movements of the eye). I think this is what JGT had in mind. The perceptual function is constructed to serve control but, once constructed, is available as an indication of the degree to which that perception is present in experience.

I don't know if JGT pointed this out, but Bill P. noted that this control based approach to perceptual learning does not lead to "arbitrary" perceptual functions. The process should converge on certain types of perceptual functions. This convergence is a result of the fact that control takes place in a "real world" -- the place that scientists are trying to model-- which is both the the input to our perceptual functions and the constraint on how we can influence perception. It is a property of the real world that "straightness" is invariant with respect to translation. So, if a control system needs to perceive a constant signal as the eye is moved left to right, it will "have to" construct a perceptual function that is sensitive only to excitation of receptors that fall on a straight line (in the plane of eye movement). This means that so called "a priori" perceptions (I think that was Kant?) might not have to be "programmed into" the architecture of the perceptual functions of the nervous system. The apparent consistency of certain perceptions (across people) -- such as perceptions of straightness, distance, whatever -- may result from the fact that these kinds of perceptual functions "must be" built in order to be able to control perceptions within the constraints of the outside world.

Maybe I should read JGT after all. He's starting to seem a LOT less silly than Freud.

Best regards Rick

Date: Wed Jun 10, 1992 12:30 pm PST
Subject: Re: Little Baby model

[Martin Taylor 920610 16:00]
(Bill Powers 920522.1330 in response to Chris Love--more of the backlog)

>The sigmoid is an interesting idea -it will put a realistic nonlinearity
>into the perceptual signal. If the reference signal were derived from
>previous recordings of the perceptual signal, it too would automatically
>have the required sigmoid nonlinearity it in; when you generate it
>directly, that isn't necessary. The loop gain will simply fall off for
>large error signals, which may well help stability. The controlled >variable will be the inverse sigmoid function of the reference signal, >when error is zero. Go ahead and try the sigmoid; I'll be interest to >see what happens. You might also try one in the output function -- I >predict that it will have remarkably little effect.

The signoid is, I think, a critical feature in both input and output. On the input side, it is the sigmoid that allows successive layers of the control hierarchy to abstract different kinds of percepts, rather than having each layer providing a rotation of the same multidimensional sensory space. In neural nets, the non-linearity is essential (though it need not be signmoidal). The perceptual connections in the control hierarchy have exatly the form of a feed-forward neural network.

On the output side, the situation may be a little different. Bill is right that it will have little effect if and only if the ECSs are more or less orthogonal, and therefore unlikely to have substantial conflicts. But when you are dealing with a denser network of control systems, as we intend in the little baby, then several ECSs will be controlling for similar percepts, for almost all percepts that concern the network. This means that there is a lot of potential for conflict. Without the sigmoid, an ECS that has a large error (and perhaps is in need of reorganizing the signs of its output links) will increasingly dominate the reference signals supplied to the next lower layer. With the sigmoid, the impact of any one "rogue" ECS is limited, and, as Bill says, its (dynamic) gain goes down as it enters the saturation region. The sigmoid allows those ECSs that can control to do so, even when they conflict with ECSs at the same level that have a large error. So I think that the sigmoid becomes critical on both sides of the network when the network contains many non-orthogonal ECSs. Non-orthogonal naturally refers to the perceptual functions of the ECSs.

Martin

Date: Wed Jun 10, 1992 2:16 pm PST
Subject: Figure 1

[From Rick Marken (920610 14:50)]

Chuck Tucker asked for a copy of Figure 1 from the "Blind men" paper.
It turned out to be real easy to do it in character form so I am posting

it to the net for those of you who might want it to go with the paper.

Figure 1

```
                          _______
     d --- k.e ---> s --- | k.o | ---> r
                    ^     _______     |
                    |                 |
                    |_____  k.f ______|
```

Pretty simple, eh?

Regards   Rick

Date: Wed Jun 10, 1992 2:58 pm PST
Subject: Re: Perceiving vs. Controlling

[Martin Taylor 920610 18:30]
(Rick Marken 920610 12:30)

> philosophically, I suppose, JGT is an empiricist -- he's saying that >we LEARN to perceive (perceptual functions are not built-in at birth.)

Yes, he was an extremist in this. In my view, he went overboard in that direction. There's no reason I can see why evolution should not have built in lots of perceptual functions, especially at low levels. Many of the same functions must be useful to cats, birds, fish, and people. So why should each individual have to learn them? But JG would have none of that. From the retinal cones to the muscle fibres, everything was assembled by the individual.

>I don't know if JGT pointed this out, but Bill P. noted that this >control based approach to perceptual learning does not lead to >"arbitrary" perceptual functions. The process should converge on certain >types of perceptual functions. This convergence is a result of the fact >that control takes place in a "real world" -- the place that scientists >are trying to model-- which is both the the input to our perceptual >functions and the constraint on how we can influence perception.

I don't know whether he pointed it out, but I assume it to be true. How could it be tested, though? We are unlikely to have the means to propose "sensible" types of perceptual function that do not, in fact, occur.

>It is a property of the real world that "straightness" is invariant with >respect to translation. So, if a control system needs to perceive a >constant signal as the eye is moved left to right, it will "have to" >construct a perceptual function that is sensitive only to excitation of >receptors that fall on a straight line (in the plane of eye movement).

Unfortunately, straightness is a property of the translation of objects in a special abstract world, not of the projection of the movement of objects across the retina. The special abstract world has no friction and no gravity. In a real world, there are tests for the straightness of edges, but except in special cases, straight things or motions in the world don't project onto straight lines or even great circles on the retina. Bill was finding this out in trying to cross-map the Little Man's kinaesthetic and visual sense of the world. JG found that with distorting spectacles, the perceptions of the straightness of straight things was regained when he acted to control them, but not otherwise. The perception of straightness is indeed a subtle and non-obvious problem.

>Maybe I should read JGT after all. He's starting to seem a LOT less >silly than Freud.

If you try to read JGT, two warnings: the writing is very dense and dry, so it is tempting to give up before reading very far, and it is not cast in the HPCT framework, so you may be put off by all the references to Clark Hull and other S-R theorists. It is their influence that JG transcended, but they provided the background for his thought.

Martin

Date: Wed Jun 10, 1992 6:18 pm PST
Subject: Perceptual Learning

[From Bill Powers (920610.1900)]

Martin Taylor (920610.1830) --

>Unfortunately, straightness is a property of the translation of objects
>in a special abstract world, not of the projection of the movement of
>objects across the retina.

This brings up an interesting point about the involvement of control in learning perceptions. As you say, straight lines aren't preserved in general when they move across the retina. Therefore if we try to create an _arbitrary_ transformation, it is more likely to destroy invariances than to create them. But suppose you take hold of an object that has a straight edge and move it so it translates (in external space) ALONG THAT EDGE. There are two results: one, the image on the retina, at least in its central parts, is not altered at all, no matter what optical distortions are involved. The second in a minute.

This means that among all the ways of moving that object across the retina _by controlling the position of that object_, there is one that will in fact create a striking invariance: a total lack of change in one perceived edge of the object (and any line parallel to it), except at the ends. Therefore if you control for the appearance of the edge, and try to keep it the same, any freedom of movement that remains will eliminate all possible invariances but one. This invariant image on the retina isn't itself "straight," as an external observer would see it. But by its invariance, it DEFINES the property of straightness.

The other effect is that a kinesthetic path will, or may, be traced by the hand holding that object in a path that also is straight in external space. So the feel of straight-line motion is generated at the same time as the look of straight-line motion. In an approximate sort of way, I can see that this may be the basis for creating Cartesian perceptual maps in both the kinesthetic and visual modalities, at the same time.

Another figure that is invariant with respect to a specific way of controlling that figure is a circle. However distorted its image on the retina, and whatever the angle of view, that image will remain unchanged if the circle is rotated about an axis normal to its center. A square is invariant with respect to 90-degree rotations; an equilaterial triangle with respect to 120-degree rotations. A plane in space is invariant with respect to particular manipulations in x,y, and z.

In each case, a particular way of moving the object as part of a process of control (and particular other ways of NOT moving it) will reveal a true invariance that is tied to geometrical properties of the outside world. All that is necessary is to control for a lack of change in the appearance of the figure while manipulating it. It seems to me that there's a principle here that could be exploited by those who are investigating perceptual learning -- but only if they do it in the context of control theory. I think it's essential for this kind of learning that one have the experience of doing things to the object yet experiencing no change in the object (that is, in the particular aspect being noticed). This shows how the object is free to change without altering its appearance. The object's invariances are learned both kinesthetically and visually by finding out what aspects of vision and kinesthesia can be changed through action without changing some particular aspect of the visual object.

Re: perceptual learning, extreme position.

J.G. Taylor may not have been out of line in insisting at all perception is learned. We have to think of two aspect of perception (which, Martin, you have actually mentioned): the TYPE of perception to be learned, and the PARTICULAR VARIANT of that type.

Consider the phasic stretch reflex. The rather complex physical arrangement that allows for this reflex involves the muscle spindle, the spinal motor neuron, and the muscle. I don't think there's much chance that the EXISTENCE of the phasic component of the stretch reflex is learned in a single lifetime. If the spindle arrangement and the sensory nerve aren't there when you're born, you're out of luck. So the TYPE of perception, the first derivative of muscle length, is present because the physical means of detecting it is present.

However, you can't be born with a damping coefficient that's fixed, if only because your arms and legs are going to grow, changing their masses and moments of inertia continually. The sensitivity of this stretch perception has to be learned; there's no other way to get the damping coeffient right. You may even have to construct new cross-connections between motor nerves serving opponent muscles. So these aspects of perception can't be inherited.

With respect to the output map I ended up using for the Little Man, I count that as more or less of a failure. It would be far better if the kinesthetic positions were SENSED in a space compatible with the visual sensory space. The only reason I didn't do it that way was that I couldn't think of a method for generating the map that would allow behavior to progress smoothly from a flat map to a map with the required curvatures in it.

Initially, I had thought that after comparing the kinesthetically sensed finger position with the visually sensed position, I could use the error to make small corrections in the map and thus create, eventually, kinesthetic perceptual signals that agreed with visual ones. But the method I used created bumps in the map, which then caused instabilities in the control process, which in turn exaggerated the bumps, and so on. What was needed was some method that would create corrections all around the position of the fingertip at any moment, so that the curvatures of the map would develop on a much larger scale. All the ways of doing that, however, took a huge amount of computation (trying to imitate a parallel system with one lonely processor). Also, the correct way of doing it would require a very long time for building the map, so that bumps would smooth out. In trying to get the map to form in some reasonable time I simply set the loop gain of the correction process too high. I felt sort of pressed for time -- I wanted to get this model into print and I didn't want to make it contingent on solving the whole perceptual mapping problem. So I settled for an output map, which works but isn't elegant. And probably isn't the right model.

Best, Bill P.

Date: Thu Jun 11, 1992 7:36 am PST
Subject: Re: Perceptual Learning

[Martin Taylor 920611 11:00]
(Bill Powers 920610.1900)

I totally agree with Bill about the symmetry properties that define "straightness" and "circularity" and their accesibility to control. I think that the use of the symmetry-based invariances is indeed the way the control systems learn to perceive these properties, and many others (perhaps all the so-called perceptual constancies such as colour and size--JGT did some experiments with size constancy along these lines, but I don't remember them.)
Martin

Date: Thu Jun 11, 1992 7:49 am PST
Subject: NS and NNSs

[from Joel Judd]

To linguists and those interested in language *(if you're controlling for relevant comments to current net discussion, skip to the last paragraph)*,

A worthwhile book to look at is:

Davies, A. (1991). _The native speaker in applied linguistics_.
Edinburgh: U. of Edinburgh Press.

He makes a good discussion of relevant questions concerning the icon of the "native speaker"(NS) and its usefulness (or lack thereof) in a theory of language. He suggests at one point that the NS might be "like the healthy person in medicine (or indeed any such state of assumed perfection) where the only definition seems to be negative, a lack of malfunction" (p. x). I also believe that the preponderance of methods of relative frequency contribute to construction of ideals such as NSs and non-native speakers (NNS). One of the nice aspects of the book is the hints Davies drops about purposeful behavior. From p. ix: "...being a native speaker is only partly about naive naturalness...the point is of course that while we don't choose where we come from we do have some measure of choice of where we go. Difficult as it is, we can change identities (even the most basic ethnicity, that of gender), we can join new groups." He also emphasizes the environmental dependencies of language: "An ill-formed language view assumes that certain forms are correct, always so, and certain forms incorrect, again always so. This cannot be so; correctness if it exists depends on context..." (p.14).

I think there might be some ammunition for the seperate/together HPCT discussions of language that have occurred in the past. Davies has a problem in defining language and community, even when borrowing Saussure's terms "langue" and "parole:"

Langue then appears to be a useful attempt to label what it is that so-called native speakers have in common...but in the event a vain attempt since it remains circular and does not help us define what it is that "langue" means. (p.19)

This discussion would be helped by the ability to define language in NON-LINGUISTIC terms such as PCT provides (though of course we USE language to talk about PCT, etc.).

Finally, later on Davies tries to define language knowledge by subdividing it into four types which relate to the recent discussion and questions on the net about perceiving and controlling. I'll just mention them for those interested

1) Metalinguistic (explicit) knowledge/"core"
2) Discriminating knowledge (what "is" and "isn't" language X)
               ##COMPETENCE##
3) Communication knowledge (social conventions--pragmatics)
4) Skills knowledge (ability--proficiency)
              ###PERFORMANCE###

A relevant comment for PROGRAM ECSs: Human grammars are not like machines that are programmed for "intelligent output." They require something to "constrain" or "restrict" what they produce--they are programs, and *part of* a larger linguistic system.

IMHO, TTFN, |=8-(<-- submitter just graduated and wears glasses but has no job

Date: Thu Jun 11, 1992 8:40 am PST
Subject: Little Man

[From Hank Folson (920611)] To Bill Powers:

When one builds a physical model of a non-PCT based little man, we know that real world inaccuracies like bearing clearances and gear backlash create problems. This got me thinking about how insensitive your little man might be to inaccuracies in its inputs and outputs, because it (he?) is a control system.

What would happen if every time in your program where you calculate a tendon stretch, joint angle or whatever, the equation is multiplied by a variable to create an inaccuracy? If you set the variable to 1, the little man will behave as intended. If you set the variable to .99 or 1.01, every effort and motion will be off by at least one per cent. But if I understand what you are doing, this will not prevent little man from controlling. He may go through some funny movements (especially if the variable is greater than unity), and go slower, but he should still get there won't he?

If several rounds of calculations are made in moving little man's finger to a target, my guess is the inaccuracy the control system can handle is quite large.

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: Thu Jun 11, 1992 9:13 am PST
From: CHARLES W. TUCKER
EMS: INTERNET / MCI ID: 376-5414
MBX: N050024@univscvm.csd.scarolina.edu

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Bill's comments on his coach

Dear Dag,

Several days agon Bill informed the net that he was unable to find his posts (I think there were two) when he mentioned his high school coach. He asked if any one had a copy of them since he could not find them. I thought of you since you mention sometime ago about your filing system. If you have those posts I would certainly appreciate them by Email if you could.

Thanks much, Chuck

Charles W. Tucker (Chuck)
Department of Sociology
University of South Carolina
Columbia SC 29208
O (803) 777-3123 or 777-6730
H (803) 254-0136 or 237-9210
BITNET: N050024 AT UNIVSCVM

Date: Thu Jun 11, 1992 9:40 am PST
Subject: Little Man

[From Rick Marken (920611.0930)

Hank Folson (920611) to Bill Powers:

>What would happen if every time in your program where you calculate a >tendon stretch, joint angle or whatever, the equation is multiplied by >a variable to create an inaccuracy? If you set the variable to 1, the >little man will behave as intended. If you set the variable to .99 or >1.01, every effort and motion will be off by at least one per cent. But >if I understand what you are doing, this will not prevent little man >from controlling. He may go through some funny movements (especially if >the variable is greater than unity), and go slower, but he should still >get there won't he?

Excellent point! One of the best technical reasons for modeing organisms as input control systems is the fact that organisms are not made out of reliable, precision components. I have looked at the effect of introducing inaccuracies into the output function of a control system (the function that transforms the neural error signal into a physical effect in the environment -- like tendon stretch). Such inaccuracies act like disturbances to the controlled variable; they are cancelled out (nearly completely, with high enough gain) by the control loop. I think it would be a great idea to have the option of introducing and eliminating these inaccuracies from the little man demo. I think the behavior of the model would look, to an observer, virtually the same in both cases (unless the loop gain was too low or the inaccuracies too great). In other words, the behavior of the little man model, with the inaccuracies, would look the same as its behavior without the inaccuracies. It will also look just like the behavior of a real person (who also has inaccuracies) -- smooth and effortless. I don't think the inaccuracies (unless they are gross) will lead to any noticeable slowing of behavior or funny movements.

This is a great idea because it shows another tangible benefit of the control system organization for robots; one that roboticists can understand.

Regards Rick

Date: Thu Jun 11, 1992 11:12 am PST
Subject: INTRO TO CSGnet & PCT

[from Gary Cziko 920611.1324]

The following is an introduction to CSGnet and PCT compiled by Dag Forssell and revised by me. To keep its length reasonable, I have deleted some information provided by Dag and I propose that he and I as well as any other interested parties work up a final form at the meeting in Durango. I also plan to eventually add more information about programs and text files available via Bill Silvert's file server. I will also eventually add information about how to deal with the listserver in order to do things obtain archive files, have your mail stopped, etc.

Any comments about this introduction would be appreciated and should be sent to both Dag <4742580@mcimail.com> and me <g-cziko@uiuc.edu>.

I plan to post this introduction once a month in addition to sending it directly to new subscribers to CSGnet. Since I have now way of knowing who is using Usenet (NetNews) to access CSGnet, posting this introduction periodically is the only way of providing an introduction to CSGnet for those using Usenet (NetNews). CSGnetters not needing an introduction may nevertheless with to consult the section describing files available on the fileserver since this will be regularly updated.--Gary

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

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

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

WHY YOU MIGHT READ THE CSGnet

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

+++
OUR SUBJECT MATTER: THE CONTROL PARADIGM

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

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

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

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

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

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

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

William T. Powers, November 3, 1991
++

THE PURPOSE OF CSGnet:

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

CSGnet PARTICIPANTS

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

THE EVOLUTION OF THE CONTROL PARADIGM

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

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

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

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

HOW TO OBTAIN TEXT AND PROGRAM FILES

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

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

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

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

get biblio.pct

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

HOW TO ASK EFFECTIVE QUESTIONS

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

When you ask a question, please consider that in order to give you a good answer, a respondent will need to put your question in context. Therefore, please introduce yourself with a statement of your professional interests and background. It will be helpful if you spell out what parts of the demonstrations, introductory papers and references you have taken the time to digest and what you learned.

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

DEMONSTRATING THE PHENOMENON OF CONTROL

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

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

Please download the introductory demonstration dem1a.exe.

THE CONTROL SYSTEMS GROUP

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

LITERATURE REFERENCES

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

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

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

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

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

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

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

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

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

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

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

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

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

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

--Dag Forssell with Gary Cziko

Gary A. Cziko Telephone: (217) 333-4382
Educational Psychology FAX: (217) 244-0538
University of Illinois E-mail: g-cziko@uiuc.edu
1310 S. Sixth Street Radio: N9MJZ
210 Education Building
Champaign, Illinois 61820-6990
USA

Date: Thu Jun 11, 1992 11:29 am PST
Subject: Robustness of Little Man

[From Bill Powers (920611.1300)]

Hank Folson (920611) --

>What would happen if every time in your program where you calculate a
>tendon stretch, joint angle or whatever, the equation is multiplied by
>a variable to create an inaccuracy?

In a way, this is already being done to some extent, as the control- system calculations are done in integer arithmetic. With a maximum signal amplitude of 200, true particularly of error signals, the inherent computational error is about 1/2 percent or more. Also, time integrations are very crude -- no Runge-Kutta stuff to improve accuracy. Doesn't affect the model.

In general, if a disturbance is inserted in the perceptual side of any control system, an exactly equivalent change in behavior is introduced: a 5% disturbance of the perceptual signal (or reference signal) will alter the controlled variable by 5%. But if this control system isn't at the highest level, a higher level system (one or more) will simply experience a little disturbance and correct it. Disturbances inserted on the output side, as Rick Marken said, are simply opposed and canceled at any level.

The parameters of the Little Man are all adjustable while the model is running. The kinds of adjustment I make to tweak up performance are like changing a "10" to a "20." That's usually enough to produce an observable change, but not always. In other words, the parameters can really slop all over the place; as long as they're roughly correct, you'll get good performance. Try doing that with one of the motor-control models they're publishing nowadays. Especially if consecutive behaviors must occur for a long time. The control-system model can run indefinitely with no cumulative error, even with inaccuracies.

>If several rounds of calculations are made in moving little man's
>finger to a target, my guess is the inaccuracy the control system can
>handle is quite large.

The entire model is recalculated 100 times per simulated second (20 times per real second on my 10 Mhz AT). So a motion that requires 0.15 sec (about the fastest) entails 15 iterations of the program.

Your intuition about the "robustness" of closed-loop models is right on target.

Best, Bill P.

Date: Thu Jun 11, 1992 3:10 pm PST
From: Dag Forssell / MCI ID: 474-2580
Subject: Intro, Coach.

[From Dag Forssell (920611)] Gary Cziko (920611.1324)

What a great post on CSG! A couple of nitpicks:

1 Requires editing:

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

2 LCS II no longer has ??? pages, but 274 or so.

3 <st> in McClelland was intended as a listserver address.

4 I have demo1. It is great. Is it realistic to expect that all prospects can and will download a demo that requires computer literacy? Could the Powers/Runkel rubber band demo (text) be offered as well? I have polished the text some more for my purposes. Will share if there is interest.

Chuck Tucker (direct) / Bill Powers (920608.0100)

>Chuck Tucker asked me for a copy of my observations about Coach in high
>school. I tried to find them and couldn't (I searched through four
>months of posts and found only that I talk too much). If anyone has a

>better filing system than I do (easy to accomplish), would it send a
>copy to Chuck at n050024@univscm.bitnet? (Those are zeros, not ohs).

Chuck, your personal (direct) plea got me off my duff. I save all posts, grouped by the week. The file find utility in Xtree Gold helped me find occurences of COACH in 9206C. Here are the particulars:

Date: Fri Apr 17, 1992 2:05 pm PST
Subject: Gateway; Testing Models [From Bill Powers (920417.1100)]

Date: Sat Apr 18, 1992 10:15 am PST
Subject: Tension, Coach [Martin Taylor 920418 13:50]

Date: Sat Apr 18, 1992 2:34 pm PST
Subject: Coach; conflict [From Bill Powers (920418.1500)]

Date: Sun Apr 19, 1992 10:59 am PST
Subject: re: testing models and conflict

To: Bill Powers and other CSGnet people
From: David Goldstein
Subject: testing models and conflict
Date: 04/18/92

Date: Sun Apr 19, 1992 11:05 am PST
Subject: re: coach; babysitter [From Bill Powers (920419.0700)]

Date: Sun Apr 19, 1992 7:38 pm PST
Subject: Re: Coach; conflict [Martin Taylor 920419 23:20]

As you can see, these all fall in the third week of April. I suggest you send a message: GET CSG-L LOG9204C to: LISTSERV@VMD.CSO.UIUC.EDU
In several hours, you will get about 220kb sent to you. Happy editing.

I am not anxious to send you the files, since I have to pay postage to MCImail. It adds up. I figure we all learn more this way.

By the way, I never did thank you for your post on rubber bands, did I? I am incorporating your illustration of how one can make another person spell something into my presentation. Thanks!

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

Date: Thu Jun 11, 1992 5:53 pm PST
Subject: getting agre's thesis

People can order Agre's thesis from the ai lab by sending email to:

publications@ai.mit.edu

it costs $14, & is AI-TR 1085

Looking at my receipt, it looks like they'll take checks or postal money orders,
though you'd have to ask them about mailing costs.

Avery.Andrews@anu.edu.au

Date: Fri Jun 12, 1992 10:31 am PST
Subject: Post-doc position

----- Begin Included Message -----

Date: Thu, 11 Jun 92 13:59:53 CDT
From: Larry Mays <OPTP014@UABDPO.BITNET>
Subject: Post-doc position
To: neuromotor-control@ai.mit.edu

POSTDOCTORAL POSITION IN SYSTEMS NEUROSCIENCE

A position is available to study the neurophysiology of the subcortical pathways mediating ocular accommodation and vergence eye movements. Some experience in single-unit recording desirable. Send C.V. and letters of reference to Dr. L. Mays, Dept. of Physiological Optics, Univ. of Alabama at Birmingham, Birmingham, AL 35294. Fax (205) 934-5725 Phone (205) 934-1158 Bitnet OPTP014@UABDPO or Internet OPTP014@UABDPO.DPO.UAB.EDU UAB is an Equal Opportunity Employer.

----- End Included Message -----

Date: Fri Jun 12, 1992 10:16 am PST
Subject: standards as conventions

[From Rick Marken 920612 11:00)

Martin Taylor (920607 1710) suggests "standards" should be viewed as conventions that make it easier to cooperate.

> It seems to me that standards
>allow you to pre-empt a possibly painful random reorganization by permitting
>you to set references that are appropriate if the other behaves in a
>conventionalized way--according to standards.

I agree that there is much to be gained from conventionalized behavior. This is particularly true in the technological world were it helps enormosly to design systems that have a standard response to actions. Thus we can be pretty confident that a clockwise turn will result in the screw going "in" or the power going "on" or "increasing". What we tend to conventionalize is the feedback function, g(o), that relates out outputs to our inputs.

>If there are any "absolute" standards, they will be those that have allowed

>the social groups using them to survive and prosper.

Conventional standards (like the clockwise turn standards) can be "absolute" to the extent that we can get all objects to abide by this convention. This can be done in principle -- though it's difficult (and sometimes not desired) in practice; some people may have a need for a counter-clockwise "in" device. But the goal of absolute standards (conventions) is at least feasible for inanimate objects because these objects have no purposes of their own that might conflict with the convention. Such is not the case with living systems.

>We can't do without standards in a time-limited social world.

Indeed, people adopt (and require) conventional ways acting for the same reason that they produce artifacts that have conventional ways of acting; it makes control simpler -- with people, it is probably what makes cooperative control feasible at all (which I think is the point of your sentence above).

The problem is that people are not inanimate objects -- and certain individuals in certain circumstances may find that acting according to a particular convention is impossible -- not because the person is bad or contrary or immoral -- but because he or she is a hierarchical control system that simply cannot act like the knob on a radio. So my argument against "absolute" standards applies as much to standards as social conventions (like grammars) as it does to standards as moral principles.
I am all for standards as conventions. I'm just saying that the notion of absolute standards -- no matter how technologically and socially helpful their existence might be -- are simply inconsistent with human nature (if people are hierarchically organized perceptual control systems). This does not mean that I believe everybody should just go ff and do their own thing. I'm just saying that this fact about human nature must be taken into consideration when we think about how people can act cooperatively.

The people who want there to be absolute standards are not "bad" people (from my point of view). The desire for absolutes is quite reasonable -- they are like people who say that "the lights should always go on when I flip the switch 'up' and off when I switch the switch 'down'". I can understand their desire -- especially with respect to people; people should never kill each other or end a sentence with a preposition; people want predictability. All I'm saying is that people are not switches; they cannot abide by such absolute conventions even if they try.
This does not mean that social chaos is inevitable (homefully, a bit less than what we now have); what I think it means is that we have to find ways to cooperate that take into account the true nature of human control systems. The fact that cooperation is possible in the context of this reality (the inability of control systems to control relative to absolute conventions) is evidenced (I think) by the general spirit of cooperation found despite the diversity (in terms of many conventions) among members of the Control System Group itself. It can be done.

Best regards Rick

Date: Fri Jun 12, 1992 10:29 am PST
Subject: Off net comment on perceptual learning

cunningB@monroe-emh1.army.mil :Bill Cunningham 920612.1345:

Forwarding a note from John Gabriel, to whom I relay much CSG-L traffic. Somebody else's attempt to reproducibly model unreproducible (portrayal of) experience was the reason for us meeting some years ago and motivates our current exploration, for which HPCT provides the paradigm./wbc

Date: Thu, 11 Jun 92 21:29:54 CDT
From: gabriel@eid.anl.gov (John Gabriel)
Subject: Re: Perceptual Learning

Yes, the symmetry theory is one of my special interests, so thanks. What Martin is saying is that the perceptual mechanism knows how things look after being turned over, rotated or whatever from experience with various objects. Circular means "same after rotation by any angle about normal to plane of object". My background says a little more than Bill P. or Martin perceive - that "well understood" experiences are all invariants, and that symmetry or invariance is an important part of learning - particularly time invariance, i.e. repeated experience from which PCT references are constructed. In fact my other serious research interest at the moment is a proof that any computer program (i.e. reproducible computation) is completely defined by one input/output pair of values, and the set of all the invariants of the program. Obvious in a way, but not observed until my recent paper to be published in Beijing. The mathematics of PROPER proof is formidable. If we take the view that any reproducible (understood?) experience can be modelled by a computer simulation, this is a perhaps a useful result about PCT.

Please forward to CSG NET if you feel this would be a useful contribution.

::John

Date: Fri Jun 12, 1992 12:27 pm PST
Subject: Perceptual learning & invariants

[From Bill Powers (920612.1300) --

I'm sending this reply to a direct post to the net (and to John) because the subject will be of general interest... here's John's post:

Date: Thu, 11 Jun 92 21:29:54 CDT
From: gabriel@eid.anl.gov (John Gabriel)
Subject: Re: Perceptual Learning

Yes, the symmetry theory is one of my special interests, so thanks. What Martin is saying is that the perceptual mechanism knows how things look after being turned over, rotated or whatever from experience with various objects. Circular means "same after rotation by any angle about normal to plane of object". My background says a little more than Bill P. or Martin perceive - that "well understood" experiences are all invariants, and that symmetry or invariance is an important part of learning - particularly time invariance, i.e. repeated experience from which PCT references are

constructed. In fact my other serious research interest at the moment is a proof that any computer program (i.e. reproducible computation) is completely defined by one input/output pair of values, and the set of all the invariants of the program. Obvious in a way, but not observed until my recent paper to be published in Beijing. The mathematics of PROPER proof is formidable. If we take the view that any reproducible (understood?) experience can be modelled by a computer simulation, this is a perhaps a useful result about PCT.

Please forward to CSG NET if you feel this would be a useful contribution.
::John

John Gabriel (920611) via Bill Cunningham (920612.1345) --

Hello, John!

If a perceptual function sees something as "the same" when it's turned over, I don't think it reports it as "the same." It simply goes on reporting it without any change. OTHER perceptual functions might see changes: apparent size, velocity, width-to-height ratio, and so on. If a perceptual function truly perceives a shape as invariant with respect to orientation, then as the shape is reoriented the perceptual signal from that function simply remains constant -- the function doesn't see any change.

But if a perceptual signal doesn't change, where is the need to control it? This brings up a rather odd feature of perceptual invariants. An invariant would remain constant under transformations that leave it unchanged. But this implies that there are other transformations that will change it. For example, if you have a square made of four toothpicks held together with clay blobs at the corners, you'll see it as square in any orientation. However, if you push in one corner, or someone does, the amount of squareness perceived will decrease. You can now control the amount of squareness perceived by pushing on other corners or pulling the disturbed one out again. So CONTROL of an "invariant" depends on applying transformations that DON'T preserve the invariance!

We tend to think of invariants (in terms of categories) as things that either exist or don't. In terms of perception, however, they simply define some canonical KIND of perception; they don't imply that this perception is unchangeable. Every perceptual function that emits a signal that is a function of multiple lower-level signals defines an invariant, in that there are ways of changing the lower-level signals that will leave the value of the function unchanged. Every weighted sum, for example, defines an invariant. If p2a = A*p1a + B*p1b, then for every way that p1a can change, there is a way that p1b can change that will keep p2a the same. But if p1b DOESN't change in that particular way, then p2a WILL change. If p2a is to be controllable by varying p1a and p1b, then the controlling system must NOT vary p1a and p1b so that A*p1a + B*p1b = constant. If it did, there would be no effect on p2a. So I repeat, control relies on causing transformations that DON't preserve the invariance. You can't "control" a circle by rotating it about an axis perpendicular to its center.

I said most of this in BCP, pp. 123, 125-6, and elsewhere, but such things tend to get lost.

Some of the confusion about invariances (and control of perception in general) probably comes from mixing control of a PARTICULAR perception with control that entails CHANGING FROM ONE KIND OF PERCEPTION TO ANOTHER. The latter case is quite different, and implies higher-level switching from one lower-level system to another as the means of control. That probably happens, but it's different from controlling the state of a single kind of perception. I don't mean that you're confused about this -- it's just something I've noticed in general. I suspect that you and I see these matters in pretty much the same way.

Once a symmetry has led to developing a particular perceptual function through learning to keep a perception constant, the next thing is to learn different states of the same perception -- that is, you perceive a circle, and then learn that you can squash it and reproduce that appearance, too, through operations that don't preserve radial symmetry. This isn't the same as perceiving an ellipse; it means perceiving a circle that's in a different state from the one you learned first. You don't see an ellipse, you see a squashed circle. The same control system that controls for circularity can also control for certain kinds of noncircularity if it's given a reference signal different from the one that's matched by a perfect circle. When presented with inputs that don't exemplify the "canonical" symmetry, a perceptual function doesn't just switch off. It reports less of the perception. And that amount of the perception, too, can become a reference signal.

<If we take the view that any reproducible (understood?) experience can
>be modelled by a computer simulation, this is a perhaps a useful result
>about PCT.

What do you mean, perhaps? If you can define what is necessary to produce a simulation of any kind of experience, this would open the door to modeling a lot of things we don't know how to model now. Keep us informed.

Best, Bill P.

Date: Fri Jun 12, 1992 1:44 pm PST
Subject: Re: standards as conventions

[Martin Taylor 920612 17:00]
(Rick Marken 920612 11:00)

>The problem is that people are not inanimate objects -- and certain
>individuals in certain circumstances may find that acting according
>to a particular convention is impossible -- not because the person
>is bad or contrary or immoral -- but because he or she is a hierarchical
>control system that simply cannot act like the knob on a radio. So
>my argument against "absolute" standards applies as much to standards as
>social conventions (like grammars) as it does to standards as moral principles.
>I am all for standards as conventions. I'm just saying that the notion
>of absolute standards -- no matter how technologically and socially
>helpful their existence might be -- are simply inconsistent with
>human nature (if people are hierarchically organized perceptual

>control systems).

I wonder if anyone was equating "standard" with "reference value". I wasn't, and if my posting gave that impression, I'll try to correct it. What I intended was to suggest that a "standard" provides a convenient level at which a reference value can be set, one that has often been found (perhaps by other people over history) to result in a desirable percept. In function, it's like the memory of a percept within an ECS.

But even with "absolute standards", there's no compulsion on anyone actually to use them as reference values. As Rick says, such use may conflict with the ability to achieve other reference values. Some day, you may have to try to kill someone if you are to maintain other desired percepts, such as personal survival or frredom.

The existence of absolute standards depends on whether over evolutionary time certain behaviours (in the PCT sense) have benefited the survival and gene-propagation of the people (or others) using those behaviours. If they have, then either by gene transmission or by social transmission, the ordinarily effective behaviours will result in absolute standards. (On social transmission, see R. Boyd and P.J.Richardson, Culture and the evolutionary process, U of Chicago Press, 1985).

I find no moral connotation to the idea of "standard", whether absolute or not.
The idea of "absolute standard" as "you have to do what I say is right" is, I think, morally and practically repugnant, for many of the reasons adduced by Rick. But "absolute standard" as "that's what people have learned as a usually effective way to behave" is simply a practical concept that improves social interaction.

Martin

Date: Fri Jun 12, 1992 2:58 pm PST
Subject: Re: standards as conventions

[From Rick Marken (920612.1330)]

Martin Taylor (920612 17:00) says:

>I wonder if anyone was equating "standard" with "reference value".

I thought you were proposing that "standards" be understood as a convention for behavior. For example, there is a convention in your country and mine that we drive on the right. So when I am on a road I try to keep my car in a lane to the right of the center line. This means (from a PCT point of view) that I set my reference for the relationship between car and center line at "right of" rather than some other value -- like "on" or "left of" of "perpendicular to". I was agreeing that standards of this sort are quite useful for successful social interaction.

>intended was to suggest that a "standard" provides a convenient level at
>which a reference value can be set, one that has often been found (perhaps

>by other people over history) to result in a desirable percept.

Well, I agree, except that I think many of these standards are fairly
arbitrary (like which side you drive on) -- they work as long as there
is agreement among those that need to abide by them in order to aviod
interpersonal conflict.

>But even with "absolute standards", there's no compulsion on anyone actually
>to use them as reference values.

Well, there is some social coersion. People can have unpleasant run ins with
the LAPD out here if they pick the wrong side to drive on. Of course, one
is still under no compulsion to set their reference at the conventional level
since they are the one's setting it.

>The existence of absolute standards depends on whether over evolutionary time
>certain behaviours (in the PCT sense) have benefited the survival and
>gene-propagation of the people (or others) using those behaviours. If they
>have, then either by gene transmission or by social transmission, the
>ordinarily effective behaviours will result in absolute standards.

If by "behaviors" you mean "references for certain inputs" then I agree;
there may be absolute (fixed, built into the individual, unvarying) references
for certain inputs. Such references are almost certainly at the cellular,
if not the genetic, level -- and they are what PCT fans would call "intrinsic
references". If, however, by "behaviors" you mean particular actions, then
I don't see how this can be correct; evolution could not possibly select
for actions that would have to produce their effects in a disturbance prone
environment. I think a lot of socio biology sounds like it imagines that
certain behaviors (in terms of actions) can evolve; for example, they
talk about evolution of "aggression". It sounds like they are talking about
the evolution of a certain visible patterns of outputs. I think the only
thing that might be able to evolve is a preference for a certain level of
sensory input resulting from these (and/or other) actions.

>   But "absolute standard" as "that's what people have learned as a
>usually effective way to behave" is simply a practical concept that improves
> social interaction.

It sounds like you are saying that an "absolute standard" is only relatively
absolute (it is usually effective at improving social interaction, but not
always). If this is what is usually meant by "absolute standard" then it
turns out that I have been advocating a version of this approach to "absolute
standards" all along. I've just been saying that some standards are usually
effective for lots of people -- but not always (they don't work for some
of the people some of the time). I just wish some of the others
in the discussion of absolute standards would have clarified this point
for me. Does this mean that the 10 commandments are "absolute standards" in
your sense of absolute standards -- it is usually effective to not steal,
but not always? Is that what judeo christians think god meant? What
about that first one -- thou shalt have no other god before me -- usually?
Some people got stewed for not obeying that one. Are some standards more
absolute than others?

No wonder we need people like Dan Qualye to help us out with this stuff.

It gets really complicated.

Best regards Rick

Date: Sat Jun 13, 1992 4:36 am PST
From: CHARLES W. TUCKER
EMS: INTERNET / MCI ID: 376-5414
MBX: N050024@univscvm.csd.scarolina.edu

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Rubber Band and Coach

Dear Dag,

Thanks for the citations; I will send for the log.

Your comments on my RB revision made me think that I had not thought about having someone spell a word - that is a great suggestion - I will incorporate it into my revision also (it adds another level of reference to the exercise). Another thought - how about spelling the word backward after you have had the other spell it forward (this reminds me of Gary's handwriting exercise). Try that one.

Regards, Chuck

Date: Sun Jun 14, 1992 8:46 am PST
Subject: From our friend at the APA

From Greg Williams (920614)

Yesterday the CSG Book Publishing office received a personal letter from the American Psychological Association which states, among other things: "Control System Group's publications are vital to the psychology community..." Wow! A foot in the door!!

The letter is signed by Bob Vrooman, Advertising Sales Representative. An ad rate card for APA journals is enclosed.

Oh, well....

Date: Sun Jun 14, 1992 10:31 am PST
Subject: Misc.

from Ed Ford (920614:11:30)

Gary Cziko -

Would you please send me your INTRODUCTION TO THE CONTROL SYSTEMS GROUP NETWORK (CSGnet). Somehow, I lost part of it and could use it as a handout when people show an interest in the CSG or in PCT. Thanks.

Martin -

Thanks for the input.

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

Date: Sun Jun 14, 1992 10:54 am PST
Subject: values, addictions

[From Rick Marken (920614 11:40)]

Well, it's pretty quiet out there. This makes it difficult for me to avoid the drudgery of writing up a paper on the "conflict-based stabilization" (to borrow Bill Powers' name for it) of improving performance in tracking tasks. So as a pass at trying to create a diversion let me try to start a thread or two. Actually, I would love to restart the standards thread because it seems like standards (or "family values") are going to play a big part in the US election. The Republicans seem to be planning to push the message that the problems we have in our society are the result of people having adopted bad values. It looks like the media is going to take the rap for having failed to teach good values and at the same time providing lots of bad value models.

Leaving aside the question of what might constitute a good or bad value -- ie. assuming everybody could agree on a set of good values, say -- I'm wondering what the Republican program is that will lead people to adopt all these good values. Is the idea to allow the media to show only material that is a model of good values (or that shows the bad consequences of having bad values)? What about the people who watch little or no TV? Movies?

Based on what I've heard the Republicans saying about the importance of family values, I think what they must have in mind is a program like this:

1) first make a list of the actual good family values
2) then publish them everywhere (especially in poor communities where, the Republicans claim, that they are most needed) and
3) accompany the published list with exhortations to "have these values, have these values"

This reminds me of the time of the French revolution when people in power had a more tangible idea of what the scruvy masses really needed and suggested the solution to them in public -- "Eat cake, eat cake!!". I sure hope the scurvy masses in the US like family values better than the french masses liked cake.

Next possible thread -- let's talk about addiction from a PCT perspective. I think this is a big problem for many people -- who feel that they are addicted to sex, drugs, alcohol (a drug), work, food, etc. Addiction of one kind or another seems to be a big problem for many people in our society.

So what is addiction from a PCT perspective? Are all addictions based on the same principles? Some addictions (like addiction to posting on CSG net) seem easier to deal with than others (like addiction to nicotine). I've been addicted to both -- I am still addicted to CSGnet (the easier addiction) but haven't smoked for 10 yrs. I suspect that quitting CSG net would be a hell of a lot easier than was quitting smoking. What's the difference? What creates a physiological addiction? Are some people more suceptable to certain addictions than others (for example, are there really "alcoholics", "anorexics", "nymphomaniacs"?). It seems like much of conventional wisdom about addiction is based on the idea that there are these "types" of people-- they have a "disease" that makes it difficult deal with alcohol, food or sex in anything but a compulsive manner. Surely PCT would say that this is ridiculous. Yet it does seem that some people are simply unable to deal with certain substances or situations in a non-addictive/compulsive manner. Why? And how might PCT help such addictive types (assuming that the person does not want to be addicted)?

Just wondering. Now, for a brief session of CSGnet withdrawal.

Best regards   Rick

Date: Sun Jun 14, 1992 11:31 am PST
Subject: Standards

[From Bill Powers (920614)]

Martin Taylor, Rick Marken, Ed Ford, whoever else is into standards --

Just a few ideas to add to the standards discussion.

Any given standard, such as "helping the poor," has at least five aspects:

1. The verbal description or name of the standard ("Helping the poor").

2. The perceptual meaning of the description or name of the standard: that is, how you can tell when a poor person is being "helped."

3. The reference level for the standard: that is, what degree of the helping is the desireable degree.

4. The program of actions used to achieve the standard: that is, what actions will help the poor to the desired degree.

5. The system concepts exemplified by the standards: that is, the concept of human nature and of society that defines the goal achieved by helping the poor.

Most discourses on standards focus on the verbal description or name of the standard, under the (incorrect) assumption that it indicates the same principle to everyone. So when old-style Democrats speak of helping the poor, they mean giving them money, advice, and services that the poor people can't obtain for themselves. When Republicans speak of the same thing, they mean doing something that will eliminate the need for giving things to poor people -- enabling them to get what they need for themselves, teaching "self-reliance."

The Republicans quite rightly claim that simply giving things to poor people will keep them dependent and poor (they don't learn how to control their own lives). The Democrats quite rightly point out that simply demanding self-reliance ends up punishing people for being poor and creates callousness toward human suffering. Republicans assume that people work in order to maintain a viable economic system that's essential to everyone, and because of financial rewards and incentives. They assume that the healthy society is one in which the members compete for wealth and predominant positions or power. Democrats assume that people work to improve the quality of their lives outside the economic framework, and that the healthy society is one in which nobody has to labor overly long, under unpleasant or dangerous conditions, or in a state of social inferiority. At least that is my view of the "canonical" positions of the two parties. I speak, of course, as a time traveler from a different era.

It's impossible to agree on standards without agreeing on system concepts: the kind of society we live in and our own human natures. Simply hurling the names of standards back and forth and claiming that they are good gets us nowhere. Even agreeing on the means of achieving standards requires a shared concept of human nature. Those who enjoy power and wealth quite rightly appreciate the advantages of these things; they advocate principles based on the assumption that everyone would be better off with power and wealth, and principles that will help those who already have power and wealth to keep them. Those who value other goals assume correctly that nobody is permanently better off with power and wealth unless everybody has them, and favor principles that spread the wealth even at the expense of those who lose out by accepting the principles.

When we speak of standards as shared principles, we tend to forget how little of this sharing there really is. The story of standards in human societies is a story of conflict, not sharing. This is true is all sizes of groups from the dating couple through the family through a whole country. Even when people say, in words, that they agree on a standard, they perceive it differently; even when they perceive it in more or less the same way, they differ on the reference level. We can agree that many poor people need immediate financial aid. Which people? How much? To be spent how? In whose Congressional district?

The other comment I have is more general. We tend to speak of standards in terms of their effects when they are shared, in terms of their roles as characteristics of a society, or in terms of what they do for social interaction. From the theoretical point of view, however, the questions are not just WHAT standards are adopted and WHY they are adopted, but what a standard IS, and HOW it can have any effect.

How does a standard influence the behavior of any individual? How does it get communicated? What has to happen inside an individual before the words describing a standard come to have meaning to that person? And what has to happen inside the person in order for any particular interpretation of such a description to attain the force of a reference condition? Without these processes internal to the individual, no standard can have either meaning or effect. We have to understand standards as they exist in and operate in a single person before we can understand how they work in a world populated by many such persons.

Finally, we often speak of the advantages or influences that standards have in a society. I think that very often, these advantages or influences are hypothetical -- they're what SHOULD occur. But I doubt that such things very often DO occur.

Best, Bill P.

Date: Sun Jun 14, 1992 11:48 am PST
Subject: Our friend at the APA

[From Bill Powers (920614b)]

Greg Williams (920614) --

>Yesterday the CSG Book Publishing office received a personal letter
>from the American Psychological Association which states, among other
>things: "Control System Group's publications are vital to the
>psychology community..." Wow! A foot in the door!!

>The letter is signed by Bob Vrooman, Advertising Sales Representative.
>An ad rate card for APA journals is enclosed.

I advocate a straight-faced reply something like this:

Dear APA,

We are delighted to hear that our publications are vital to the psychological community, but somewhat puzzled because said community has paid little or no attention to the work of the CSG. Could you explain just what it is about the publications of the CSG that caught your attention and elicited this flattering statement?

Best, Bill P.

Date: Sun Jun 14, 1992 12:22 pm PST
Subject: Blind men

[From Rick Marken (920614.1315)]

Gary Cziko just told me he did not receive a copy of the blind men paper. I think I have been having some problems sending personal e-mail -- but with no feedback I have no way of knowing. So, if you asked for a paper and didn't get it, please feel free to ask again. I sent several but have no way of knowing if the sending was successful.

Regards Rick

Date: Sun Jun 14, 1992 1:42 pm PST
Subject: Butler on addiction AND standards

From Greg Williams (920614 - 2)

Samuel Butler (the Erewhon guy, not the poet) long ago had a comment related to both the proposed thread on addiction and the one on standards:

If the hangover came first, alcoholism would be a virtue.

An aside to Bill Powers: are you addicted to revisions? Let me know when you are getting tired of revising, and THEN I'll have a go at the paper. In the meantime, I have made some literature-survey progress.

Best wishes (assuming that meets standards and isn't too addictive),

Greg

Date: Sun Jun 14, 1992 4:31 pm PST
Subject: standards

[From Rick Marken (920614.1700)] Bill Powers says:

>It's impossible to agree on standards without agreeing on system concepts:
>the kind of society we live in and our own human natures. Simply hurling
>the names of standards back and forth and claiming that they are good gets
>us nowhere.

I guess I understood this to be true even before I learned PCT. Alot of PCT is just common sense -- although this is apparently not true for certain species of Quayle.

>When we speak of standards as shared principles, we tend to forget how
>little of this sharing there really is. The story of standards in human
>societies is a story of conflict, not sharing.

How true, how true.

>Even when people say, in words, that they agree on a standard, they
>perceive it differently;

Language has certainly been a mixed blessing, no? And now with the dominance of "symbolic" approaches to cognition it is almost impossible to get people to pay attention to their perceptions (and imaginations) instead of the words that are used to describe (and point to) them.

>Finally, we often speak of the advantages or influences that standards have
>in a society. I think that very often, these advantages or influences are
>hypothetical -- they're what SHOULD occur. But I doubt that such things
>very often DO occur.

This has been my point all along -- at least in terms of personal (and, to some extent in terms of interpersonal) control. I see no way in which perceiving certain standards at certain reference levels can NECESSARILY lead to successful control of any other perceptual variables -- or intrinsic variables. Yet this is an article of faith for many people in society. I imagine that if a PCTer showed, quantitatively and experimentally, that this faith is not correct he or she would soon be the victim of a holy war. Now that I think of it, if real good science, on control of principles and system concepts were done, it is possible that the results would make the religious/political/scientific establishment take steps that would make the Cathlolic church's treatment of Galileo look like a picnic.

I think I'll just do the manual control stuff (no variables above level 8 (programs?)).

Regards Rick

Date: Mon Jun 15, 1992 6:12 am PST
Subject: causation

An idea I thought I might mention concerning causation and behavior. Between comments made on the net and some thinking of my own, I have found two means of expressing causal relationships concerning volitional behavior. Which of the following sets of statements is preferable? I find #1 to be more precise in meaning, #2 to be more to the point.

1) Outcomes are influenced by the environment and reference signals.
   Outcomes are determined by reference signals.
   Actions are influenced by the environment and reference signals.
   Actions are determined by the environoment.
   Actions control outcomes.

2) The cause of an outcome is the reference signal with the
   environment as the enabling condition.
   The cause of and acton is the environment, with the reference signal
   as the enabling condition.
   Actions control outcomes.

I can add that reference signals within the hierarchy, when the hierarchy is viewed as a whole, are determined #1("caused, #2) by the environment since they are the integrations (summations, combinations?) of higher level outputs, which as I said above #1 are determined by the environment. Of course if we are refering to a reference signal at, say, level 4 and we begin our analysis at level 4, then we do not speak as if that reference signal is determined by the environment since its a derivative of the "autonomy from the top." But if we took a step back and viewed the entire hierarchy, we would say that the reference signals at level 4 are determined by the environment.

I would appreciate if replies could be sent to my mailbox at m-olson@uiuc.edu since I can't keep up with the net's pace. See ya in Durango!

Carpe' Diem Mark Olson

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 Jun 15, 1992 9:21 am PST
Subject: Causation

[From Bill Powers (920615.0900)]

Mark Olson (920614.1700) --

>1) Outcomes are influenced by the environment and reference signals.
> Outcomes are determined by reference signals.
> Actions are influenced by the environment and reference signals.
> Actions are determined by the environment.
> Actions control outcomes.
>
>2) The cause of an outcome is the reference signal with the
> environment as the enabling condition.
> The cause of an action is the environment, with the reference
> signal as the enabling condition.
> Actions control outcomes.

This kind of summing-up is a good idea, I think. Let me throw in some more considerations for you to sort through.

I think "the environment" is perhaps too general, because properties of the environment enter differently from independent variables in the environment. A lever has the property that pushing down on one end makes the other end go up. So it determines that IF you want the far end to go up, THEN you push down on the near end. Skinner used the term "contingency" for this: movement of the far end is contingent on movement of the near end. Not a bad term.

But the lever neither causes nor influences anything by just sitting there. The variable position of its far end, given the lever's properties, is influenced in a particular way by the action of moving the near end. It is also influenced by independent forces acting anywhere along the lever. The lever doesn't determine whether any such actions or forces will occur.

Maybe this could be summed up by saying that the KIND of outcome of an action is determined by properties of the environment, while the STATE of the outcome is influenced by action and independent variables: this comes down to JOINT DETERMINATION of the state of the outcome by action and disturbance, given the properties of the feedback connection.

Of course this doesn't bring purpose into it yet. There are always joint outcomes of action and disturbance, but INTENDED outcomes are a subset: those maintained near a given state by variations in action.

Shoot, now I'm getting into details that seem to go on forever. For example, what is an "outcome?" Isn't it defined by perception? When you act on an environment you cause a never-ending series of shifts in relationships among thousands of detailed variables. It's perception that puts order into these variables, represents them as abstract things like "objects." Outcomes are defined in terms of things we perceive. While perception doesn't CAUSE the physical variables to be in certain states, it SELECTS certain functions of those variables as significant aspects of the world, and intended outcomes are among those aspects.

Oh, blah, blah, blah. Suddenly I'm wondering if there really is any way to reduce the relationships to a sequence of simple statements of influence, determination, or causation. Those concepts are all really two-element ideas:

the influence or determiner or causer, and the influenced, determined, or caused. Maybe all this boils down to trying to express circular causation in terms of lineal causation. I'm beginning to see why Aristotle ended up with four kinds of causes.

I guess the problem with your summing-up is that it makes perfect sense if you already understand the principles of control, but that each statement requires endless explanation for someone who doesn't. In a real control system, there are all kinds of determinations and causations and influences going on at the same time. Perception defines what amounts to an outcome. The environment, given that definition, determines what actions would be needed to affect that outcome in any given direction, and also determines how any remote independent variables that may exist can affect that same outcome. Reference signals specify which state of a perceived outcome is to be maintained; disturbances that cause errors relative to that reference state then determine what amount and direction of action, applied through properties of the environment, will occur.

I feel the same way you do: there has to be a way to express all this as a series of easily-understood statements. But after all the above, I'm beginning to wonder whether ordinary language doesn't contain too much SR theory for this to be done using existing words. Just talking about causes and effects is a problem, and influences and determinations really don't help that much. Maybe in a subtle way we're trying to make control theory palatable to people who still think in straight lines.

After all that, I guess I just should have said "Very interesting." But what you're trying to do deserves more than that.

Best, Bill P.
Date: Mon Jun 15, 1992 1:21 pm PST

[From Dennis Delprato (920615)]

Rick Marken (920614) refers to Republicans and Democrats, the leaders. I recently heard the Presidential candidate of the Libertarian Party quote another (foreign) leader. To paraphrase:

> You Americans attempt entirely too much to control your
> citizens. It won't work. I know for I tried.

According to my source, this came from none other than Hermann Goering--at his trial in Israel.

This sort of observation brings home just how far removed from our world as lived is PCT. Tied in with this is the fact that despite widespread dissatisfaction with Republication and Democrat policies, few people seem to have any idea of what Libertarians are all about.

The moral here may be to stay quite close to pure science for the time being. I assume many members of this network have thought about the Libertarian-PCT connection.

Date: Mon Jun 15, 1992 2:20 pm PST
Subject: Correction

[From: Dennis Delprato]

Check assertion in earlier message today that Herr Goering made it to Israel. Obviously I need to watch more movies--such as Judgment at Nuremberg. That is where, of course, he performed the Supreme Act of Self-Control that some classify as escape behavior in his particular case.

Date: Mon Jun 15, 1992 3:22 pm PST
Subject: Politics and PCT

[From Bill Powers (920615.1630)]

Dennis Delprato (920615) --

>This sort of observation brings home just how far removed from our world as
>lived is PCT. Tied in with this is the fact that despite widespread
>dissatisfaction with Republication and Democrat policies, few people seem to
>have any idea of what Libertarians are all about.

Hmm, if that first sentence is true, then PCT needs a lot of work, doesn't it? Perhaps some amplification on these failings of the theory is in order? Although I'm glad to hear your voice on the net again, you bear alarming news!

I've heard of the Libertarians, as I suppose most people have. My impression, based on very little contact with them, is that they want to be completely free of coercion by government or other people. This is understandable from the standpoint of PCT; no person wants to be subject to coercion by others, either single or collective others. Everyone wants laws that will keep others from being dangerous or inconvenient to oneself; everyone wants to avoid taking on burdens from other people; everyone wants to keep what belongs to him or her and work only as much as necessary to get it in the first place.

Of course PCT also explains why people DO attempt to control others and why this normally leads to coercion, violence, and all that as the controllers try to get control and the rest resist it.

I think people want other things, too: that is, we aren't all really thinking only of ourselves. I have notions about the kind of society I'd like to live in, and to some extent am willing to work harder than actually necessary to keep myself alive and happy in order to feel that I'm furthering that sort of world. I'm willing to and free to abide by certain collective agreements that seem best for all even though they're sometimes inconvenient for me (really, if I forgot my money, why can't I just walk out of the restaurant without paying? What's the big deal?).

I think it would be interesting to discuss political positions from the standpoint of PCT. This would be, of course, like discussing religion -- teetering on the edge of offending everyone. But if we could maintain a sort of anthropological distance from the subject, we might learn something about it. Every political stance has good points and bad points -- I'm a fan of the sometimes very sensible American Socialists, for instance, although I don't vote that way. I've always thought the ideals of Communism were pretty attractive, although had I ever traveled to the old USSR I would have instantly been clapped into a Gulag because I can't keep my mouth shut. I can even see some attractive points in right-wing conservatism. In other words, all political movements have some justification for existence that makes sense to someone, and that ought to make sense in terms of PCT. The reason the whole world doesn't have just one political movement is that there's never been one that deals competently with ALL the problems people want solved. Maybe there can't be, for the reason that people subscribe to conflicting system concepts -- for example, Democrats, Republicans, Libertarians, Socialists, Communists, and so on.

>The moral here may be to stay quite close to pure science
>for the time being. I assume many members of this network
>have thought about the Libertarian-PCT connection.

If pure science has to avoid politics, it's not good for much in the real world, is it? This would be like confining HPCT to the first 9 levels and declaring the last, or rather next, two out of bounds.

This seems to be a good season for this subject -- shall we see if we can manage it without ending up enemies forever? How about kicking it off, Dennis, by describing the system concept that the Libertarians are proposing? Anyone else who wishes can play devil's advocate for other political points of view. I suggest that we confine ourselves to describing and justifying the reference levels and perceptions, and at least at first not be concerned with contrasting them with those of other systems. What is the best of all possible Libertarian (Democratic, Republican, etc) worlds and why is it the best?

Best Bill P.

Date: Tue Jun 16, 1992 6:27 am PST
Subject: Politics & PCT

[From Dennis Delprato (920616)]

>Bill Powers (920615.1630)

>>This sort of observation brings home just how far removed
>>from our world as lived is PCT.

>Hmm, if that first sentence is true, then PCT needs a lot of work, >doesn't it? Perhaps some amplification on these failings of the theory >is in order? Although I'm glad to hear your voice on the net again, you >bear alarming news!

No failings or flaws are implied. PCT likely _would_ be flawed if it reiterated what is already embraced by the mainstream. By "the world as lived" I mean how things are, what is currently believed, as contrasted with what remains to be discovered by those who have not looked with eyes unblinded by their past. My news thus does not seem to be very new. Witness complaints regarding acceptance of PCT.

>If pure science has to avoid politics, it's not good for much in the >real world, is it? This would be like confining HPCT to the first 9 >levels and declaring the last, or rather next, two out of bounds.

Another view here is that pure science (that is particular aspects of science such as why and how organisms function psychologically) is best developed with as little an eye on extraneous matters as possible. For example, it is possible that PCT as a fundamental theory of relevance for what today is called psychology will better progress if those few of you who really know what is going on stick with PCT in its purest sense. According to this view, the time you spend thinking, writing, and talking about politics will simply allow less time to explore PCT. This view takes applicability of the theory to politics as a low priority in the developmental phases of the theory.

>This seems to be a good season for this subject -- shall we see if we >can manage it without ending up enemies forever? How about kicking it >off, Dennis, by describing the system concept that the Libertarians are >proposing? Anyone else who wishes can play devil's advocate for other >political points of view. I suggest that we confine ourselves to >describing and justifying the reference levels and perceptions, and at >least at first not be concerned with contrasting them with those of >other systems. What is the best of all possible Libertarian (Democratic, >Republican, etc) worlds and why is it the best?

See above paragraph. On the other hand, I suppose someone could do something very interesting with political control (and notice that politics seems all about control) and PCT. Despite flaws in Skinner's operant psychology, his utopian Walden Two is much more appealing to me (at times) than a society run by the likes of Nixon, Carter, Bush, Clinton, Clarence Thomas, et al. Of course, democracy is supposed to give citizens control over their own fate, and we are a lot better off than many others on this. But there are many reasons for citizens to support rigid controls over others. For example, it is possible that I would be better off if I appealed directly to the citizens of Michigan for support, but right now I am inclined to be silent about how they are coerced in my behalf. And my case is extremely mild as compared with flagrant abuses of the authority to strong arm others. Actually, I know very little about the Libertarian Party either in historical perspective or in terms of where they are today. Libertarians seem to give individual freedom the highest priority and seem sensitive to many cases of verbal claims of looking out for the individual used to becloud coercive tactics.

I do hear occasional claims that pure conservatism above all seeks to keep the government as uninvolved in individuals' lives as possible. We detect very watered-down versions of this in various Republican-sponsored policies. As with the Libertarian position, as I understand it, the conservative would minimize disturbances to social and individual systems. Thus, the systems would require less effort directed at adjustment, permitting more harmonious functioning.

Are there any political scientists around?

Date: Tue Jun 16, 1992 6:44 am PST
Subject: Addiction

[From Hank Folson (920616)]

Some anecdotal information on addiction from a conversation on a bicycle ride with a fellow who evaluates drug programs: 1. Crack babies can fully recover, but they need a great deal of touching, holding, and hugging. 2. There are drug programs that work. We didn't discuss success rates. They are cost effective, mainly because avoiding one crack baby saves $350,000. 3. Detoxification of addicts is only the beginning of treatment. Long term addicts have no experience making decisions, and may be taking drugs to avoid making decisions. Recovering addicts need a lot of counseling time to train and support them in making decisions about their lives. 4. Conservative Republicans are against even cost effective programs as they perceive addiction as the result of a lack of will power.

I do not use drugs and my knowledge is limited to a TV special on PBS, so I am not controlling for approval of my PCT thoughts on cocaine/heroin addiction:

The mind functions by chemical action, so our reality is chemically based. When we are sober, we have one reality, call it a sober reality. When mind altering drugs are in the brain, they interact with the chemical action of the brain to create a second drug based reality. One anecdote I recall from a PBS special is that researchers were surprised at how many returning Viet Nam War veterans dropped heroin habits. In war, the sober reality can not compete with the drug reality. But once these soldiers left the Viet Nam reality, wouldn't the sober reality at home be preferred over the drug reality in Viet Nam? The addiction would disappear for those who came home to a happier environment.

Our reality is a running average of the daily experiences in our lives, weighted by the impact of various experiences. My guess is the drug reality neither adds onto nor overwrites the sober reality, but exists along side of it. The person can connect the drug reality to instances of drug use, and so keep the two realities separate. The next step in my supposition is that we compare the two realities. If the drug reality produces lower error signals than the sober reality, the user would, I think, control for the drug reality. The behavior that results is more drug use, as the user controls for what he perceives as a better reality.

The drug reality is completely inside the user's head, insulated from the outside world, so it stays much the same no matter what happens to the sober reality. The financial and social costs of supporting the habit degrade the quality of the sober reality, so the sober reality relatively looks worse and worse. The increased error signals will lead to controlling for the drug reality. The addicted behavior becomes stronger, and harder to break.

Soberly, 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: Tue Jun 16, 1992 10:39 am PST
Subject: Meeting scholarship

[From Bill Powers (920616)]

Joel Judd regrettably will not be able to attend the CSG conference and has given up his scholarship. Accordingly, there is one scholarship available for a student to attend the conference (covers conference fee, meals, and lodging, but not transportation or $5 student membership). Contact Mary Powers, 73 Ridge Place, CR 510 Durango, CO 81301 immediately if you wish to apply for it -- first application with faculty member approval takes it.

Best, Bill P.

Date: Tue Jun 16, 1992 11:45 am PST
Subject: misc catchup

[From: Bruce Nevin (Tue 920416 12:22:10)]

My my my am I behind the stream.

(Mark Olson (920614.1700)) --
(Bill Powers (920615.0900)) --

Since causation is circular, what is needed perhaps is to replace the punctuation "." with things like ", and" and ", which" and a (graphically) final ". . .":

```
    Outcomes influence or determine some other environmental variables
      and determine some reference signals
      (presumably when intrinsic error is low?);
    and some environmental variables and reference signals influence actions
      while other environmental variables determine actions;
    and actions Control outcomes
    and outcomes influence or determine other environmental variables . . .

        +---------------+       Everything in this diagram is outside
        |               |       the observed control system except R;
        v               |       that is, everything here is made up of
        E -->           |       "environmental variables," and R is a
              A --> O --+       record in memory of perceptions of
        R -->           |       environmental variables that the
        ^               |       control system maintains as desirable
        |               |       goals.
        +---------------+
```

>1) Outcomes are influenced by the environment and reference signals.
> Outcomes are determined by reference signals.
> Actions are influenced by the environment and reference signals.
> Actions are determined by the environonment.
> Actions control outcomes.

How do reference signals influence outcomes other than by way of actions?

One thing that is wrong with the above diagram is brought out in part by your second formulation, in the statement that R causes O (enabled by

E), and E causes A (enabled by R).

```
>2)  The cause of an outcome is the reference signal with the environment as
>      the enabling condition.
>    The cause of an action is the environment, with the reference signal as
>      the enabling condition.
>    Actions control outcomes.
```

It might help if you recapitulated your definitions of cause, influence, and enable.

```
     ............
     :           :
     :   +--E=R--:--> A --> O <-- E
     :   |       :          |
     :R [_]      :          |
     :   ^       :          |
     :   |       :          |
     :   +---I---:----------+
     :...........:
```

This seems to me a more accurate formulation:

```
        Comparison of sensory input I with reference R
        determines error E=R (collapsing levels)
        which determines action A
        which with other environmental variables E determines outcome O
        which determines sensory input I
        and so on.

        It is only by continuous adjustment of E=R and thence of A
        that R controls O.
```

It is the continuous nature of control that is perhaps hardest to bear in mind, that it is not stepwise sequential causation.

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

(Dennis Delprato (920615)) --
(Bill Powers (920615.1630)) --

"Libertarian" used to be the term of choice for the anarchists, before the last Depression. It has been taken over by Invisible Hand fans, like Ayn Rand's devotees. An interesting view of how anarchism works is in Ursula LeGuin's novel _The Dispossessed_. A fine writer. Her father was Alfred Kroeber, the "father" of American anthropology.

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

(Rick Marken (920614 11:40)) --
(Hank Folson (920616)) --

I thought that addictive substances grabbed you by the intrinsic error. That includes addiction to substances that the body normally produces (endorphins, adrenalin, etc.) It's easier to find substitutes for some

(CSGnet traffic) than for others (nicotine).

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

(Andy Papanicolaou & Tom Bourbon (920606)) --
(Bill Powers (920610.0800)) --

RE: habits and speech

>So, we believe that it is generally true (and, hopefully, the linguists
>will concur or persuade us otherwise) that every time a /ba/ or a /ga/
>or a /tu/ is heard the corresponding patterns of articulatory gestures
>contain a set of invariant features.

The search for invariant articulatory features for the phonemic contrasts of a language continues to be problematic. As discussed some months ago, articulatory (tactile and kinesthetic) perceptions seem more important for consonantal sounds and acoustic perceptions for vocalic sounds, and acoustic perceptions seem in general most important, but cannot be dispensed with--note for example that when one has a local anesthetic in one's mouth one's speech is not disturbed nearly so much as one is convinced it must be, but that it is disturbed (largely by pushing harder for tactile sensations that are attenuated by the anesthetic, whatever kinesthetic perceptions there may be are harder to judge).

The reason it is problematic, in my opinion, is that the actual pronunciations are a byproduct of control for several different kinds of perceptions. Primarily, one is controlling for contrast (e.g. all words that have p in a given syllable position as against all words that have b or m or t in the corresponding environment). Secondarily, one is controlling for doing so in a manner associated with some social group (community, social class, etc.) and/or in a manner different from that associated with some other social group. Also secondarily one is controlling for variation in the gain on the above two kinds of control (perhaps by definition this is tertiary control wrt to those), where the variation communicates affect, attitude toward the information communicated, relationship to conversants, self-image, etc., in ways that are as yet still poorly understood. Thank goodness articulation is relatively immune to disturbance from the environment! Though there must be some of that as well. Certainly one varies the gain depending on environmental noise, etc. The actual articulatory movements vary all over the map for different speakers, and even for the same speaker under different conditions. Peter Ladefoged I understand did some interesting experiments involving jaw movements in speech of people suspended upside down, for example.

I should think Martin would have quite a bit to say about this.

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

(Joel Judd (Thu, 11 Jun 1992 10:26:17)) --

The myth of The Native Speaker is as you say the flip side of the myth of Language (langue) as a monolithic ideal Platonic Reality. Both concepts ignore heterogeneity and variation and therefore are covertly normative. A good reason for ignoring heterogeneity and variation is that simplification is a prerequisite to first steps of progress. A necessary reason for doing so, and the reason that it is covert, is that in the nature of social injunctions of a certain order one must learn them, forget that one has learned them, and forget that one has forgotten (in R.D. Laing's way of putting it), yet nonetheless continue to obey them. Thus the crass confession will not do, "speak in this manner if you want social advantages." One's manner of speaking must be unconsciously controlled, not manipulated, if it is to be authentic. Persons who have learned to manipulate body language are if discovered profoundly distrusted by their fellows, and for good reason.

>I think there might be some ammunition for the seperate/together HPCT
>discussions of language that have occurred in the past.

Can you unpack that a bit? What do you mean "separate/together"?

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

One of the problems of individual and community is that mammals and probably other animals take their lead from their fellows in setting some reference signals. Much of how we do things with words depends on this. There is an interesting survey of primate research in a new book by Carl Sagan and his wife Ann Drayan, _Shadows of Forgotten Ancestors_. The following passage is from an excerpt published in _Parade_ magazine (for June 7) which comes with the Boston Globe on Sundays:

> The perks of being an alpha male entail certain obligations. In return for deference and respect, for preferential sexual access to ovulating females and for deluxe dining privileges, he must render services to the community, both practical and symbolic. He adopts an impressive demeanor, even something approaching pomp, in part because his subordinates demand it of him. They crave reassurance. They are natural followers. They have an irresistable need to be led.

A social and political structure that is easily recognizable as human-like, complete with palace intrigues and Marshall Dillon style enforcement of lawn order, has evolved and is maintained among chimpanzees without language as we know it. This is not to deny that we do the same with language, and more elaborately by virtue of having language. The point for me is that we use language for such purposes in parallel with the essential function of language, which is error-free transmission of information. Our use of language for social status, political process, communication of affect, and so on is essentially gestural, just like our use of facial expression, posture, and so on. The free bandwidth of facial expression, after serving more basic functions like seeing, smelling, consuming food, and the like, is used to communicate affect, relationship (including relative status), etc. Just so, after the basic function of transmitting particular dependencies of words there is free bandwidth of pronunciation, word choice, choice of paraphrase, differentiation of topic and comment, and so on that is used to communicate affect, relationship to the information, relationship to other persons, etc. Although each choice within this bandwidth is "verbal," selection among available choices is just as "nonverbal" as a smile or a stamp of the foot, and is controlled in consonance with them.

I recommend Edward Sapir's perspective on the relationship of individual and society, and for insight into it I recommend Harris's review of Sapir's _Selected Writings_ which appeared in _Language_ 27:288-333 (1951).

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

I have several times posed the question, how can an ECS controlling for perception of a word (or category) distinguish between tokens that have the same referent and those that do not. One approach that I have not mentioned is to suppose that a word (or a category) has the same reference as a repetition of it has unless they are explicitly differentiated. An occurrance of a word may be differentiated from another occurrance of the same word by asserting something of the first and denying the same of the second. By allowing for zeroed context, including dictionary sentences making explicit relations like antynomy, classification, and so on, any less obvious differentiating context expands to an assertion/denial pair.

Now all of this could be done with nonverbal perceptions, but only if there is a way of controlling the assertion/denial pairing without use of words. My understanding, and my subjective impression, is that there is no nonverbal equivalent of denial or negation. Prohibition, yes, but that is different.

So previously the question was how to capture the perception of sameness nonverbally; now it is how to capture the perception of difference between two tokens of the same perceptual signal. The alternative approach mentioned here does not seem any easier, and that is why I did not mention it before.

Any insights?

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

I want to work up some ideas on the reduction of logic and the program level to simpler terms based on G. Spencer Brown's _Laws of Form_. Gregory Bateson touted this book in his talk to the annual meeting of the American Anthropology Association in the summer of 1974. A few years back I met a person who had implemented an inference engine based on it, using a language he called LOSP. I want to find his papers before I say more. But Brown's book is available. It restates Principia Mathematica and more in terms I think very congenial to HPCT, and I know very congenial to Harrisian linguistics. It all starts with making a distinction (contrast). Everything falls out from that essential, first arbitrary act.

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

Gotta run. Bruce bn@bbn.com

Date: Tue Jun 16, 1992 2:52 pm PST
Subject: Temperature control

[From Bill Powers (920616.1600) (The Control Systems Group)

Fang Zhong (Duke University 920616) --
Copy to CSGnet

Hello, Fang Zhong. You've been put in contact with a conference concerned mainly with control theory as a tool for understanding behavior. As we use pretty simple versions of control theory, maybe you've come to the right place for a simple answer.

I have to make some guesses about the physical setup. I assume that your heater is in contact with some thermal mass, perhaps water. So the mass will heat up at a rate proportional to input power. There will also be a thermal transport lag if the sensor is at all far from the heater. It's mainly the transport lag that will cause the oscillations when you raise the sensitivity. If you're using water, you can cut down on the transport lag by madly stirring the water. But that's not your main problem.

Consider this line:

```
>   vh2 = prev_vh * prev_vh + (2 * cur_r - target_r - prev_r) * dv_dr
```

If you want to control temperature, you have to compare sensed temperature with reference or target temperature. So the error signal would just be (target_r - cur_r).

As you are using 2 * cur_r and subtracing prev_r, I assume you're adding some first derivative of the error signal: that is,

```
  total error = target - cur_r  -    (cur_r - prev_r)
                 <--error -->     <-- deriv of error -->
```

This is my standard way of writing the error signal -- the sensor signal is always subtracted from the reference signal. No particular reason, that's just my convention. It keeps the feedback negative with all other constants in the loop positive.

Now you want the heater voltage vh to be the integral of the error signal. If you add the error signal to the SQUARE of the heater voltage (vh2), you'll get a hybrid between the square of the integral and the integral of the square. I don't think you want either one. It would be best to compute first just the integral of the error:

```
   vh2 = vh2 + (target_r - cur_r) - (cur_r - prev_r)
```

Now vh2 doesn't mean the square, it's just used as a dummy variable (initialized to zero).

If you want heating rate to be linear with error, and I think this is best, you should output a heater voltage proportional to the square root of vh2, because power output goes as vh*vh. So your step

```
        vh = sqrt(vh2)
```

is OK. It would be best to precede this step by

```
 if(vh2 < 0) vh2 = 0
```

to avoid accidentally taking the square root of a negative number.

This will result in a linear heating system with integral error control, some first-derivative phase advance, and another integration in the environmental part of the loop that converts power output to rate of change of temperature (temperature is roughly the integral of power output). If the derivative contribution is large enough, you will get a system that looks like a single integrator, and it will be stable save for the effects of transport lag.

You need to be able to vary the amount of first derivative in relation to the error signal, changing the first program line above to:

```
   vh2 = vh2 + k1*(target_r - cur_r) - k2*(cur_r - prev_r).
```

By setting k1 VERY small and k2 zero, you can get a system that approaches the reference temperature and stays there. It will act very slowly. Then you can start increasing k1 to create oscillatory control, then k2 to eliminate the oscillations. Eventually you will arrive at values of k1 and k2 that give you the fastest possible control without oscillation. You won't be able to go any faster than that because of transport lags.

Let us know how it works!

Best Bill Powers powers_w%flc@vaxf.colorado.edu

Date: Tue Jun 16, 1992 7:34 pm PST
Subject: Causation, determination, and words

[From Bill Powers (920616.1930)]

Bruce Nevin (920616) and Mark Olson (920614) --

" Outcomes influence or determine some other environmental variables and determine some reference signals (presumably when intrinsic error is low?); and some environmental variables and reference signals influence actions while other environmental variables determine actions;
and actions Control outcomes and outcomes influence or determine
other environmental variables . . ." (Nevin)

This attempt to put the basic relationships of HPCT into simple language is getting more and more confusing. Let's go back to the basic diagram and try to build it up in an orderly way.

First, a simplified ECS diagram showing an Input function, a Comparator, and an Output function. In the environment we have a controlled variable or input quantity qi, an output quantity qo, an environment function E, a disturbance function D, and a disturbing quantity qd. Inside the organism we have a reference signal r and a perceptual signal p (put your own arrow heads in).

```
       p  r
       |  |
     -----------
    | I | C | O |
     -----------
      |       |
     qi<-E<- qo
      |
     D <- qd
```

Where do we start? By assuming that control is good, so the input quantity matches the reference signal.

$$(1)\ qi = I^{-1}(r)$$

The input quantity is maintained at a value which is the inverse of the input function of the reference signal. So:

(a) The reference signal and input function determine the state of the input quantity in the environment.

Statement (a) is an approximation. In fact, the reference signal determines the state of the input quantity within some region r plus or minus epsilon, where the size of epsilon depends on the loop gain of the control system and the maximum disturbance that the system can resist. Using the approximation implies assuming an ideal control systems (infinite loop gain). In that case, epsilon is zero.
Next, we explain HOW this determination is brought about. qi remains near the specified state because variations in the disturbing quantity, transformed through the disturbance function D, are opposed almost exactly by variations in the output quantity transformed by the environmental feedback function E. Therefore:

$$(2)\ E(qo) + D(qd) = I^{-1}(r), \text{ or for later reference,}$$

$$qd = E^{-1}(I^{-1}(r) - D(qd)$$

The sum of the output quantity and the disturbing quantity, each transformed by the appropriate function in the environment, must equal the value of input quantity determined by the reference signal. This leads to the statement

(b) The reference signal and external disturbances jointly determine the output quantity.

Statement (b) says that given a constant reference signal, variations in the disturbance call forth specific variations of the output quantity or action, in the manner of an apparent causal relationship. The determining effect of the disturbance on the output, however, is subject to the condition that the sum of disturbance and output effects always equal a particular value: the value of the input quantity determined by the reference signal. This balance point, therefore, can change if the reference signal changes. This is why the action of the system is JOINTLY determined by disturbances and the reference signal, and not exclusively determined by either.

Note that in both statement (a) and statement (b), the apparent causal relationship works in the opposite direction to the direction of physical causality. The reference signal appears to determine the input quantity backward through the input function. The disturbance appears to affect the output quantity (via the input quantity) backward through the environment function. The form of the apparent causal relationship is, in both cases, the inverse of the function actually connecting the variables by the most direct route. These backward relationships are a direct result of the closed-loop organization.

There is only one dyadic deterministic relationship: the reference signal determines the input quantity. The output quantity depends on two variables, jointly: the disturbance and the reference signal.

Now let's put together a two-level system:

```
         p2 r2
          |  |
        -----------
       | I2| C | O |
      / -----------
    /    |        |
  qi2a   qi2b    qo2
  |      |     /
  |      |    /
  |      p   r
  |      |   |
  |     -----------
  |    | I1| C | O |
  |     -----------
  |      |        |
qi1a   qi1b<- E<- qo1
  |      |
  D2     D1 <- qd1
  |
 qd2
```

This is a case in which the higher-level system generates a perception derived from two lower-level perceptions, one of which is controlled and one of which is not.

Now we can see that the first-level reference signal is the output of the second-level system. The reference setting at level 2 can be seen in behavior only if we look at the two input quantities qi1a and qi1b through the same kind of input function that the second-level system uses. If the second-level system perceives the sum of qi1a and qi1b, we must observe the sum of these input quantities in order to see what is being controlled. Furthermore, we must see the two disturbances, qd1 and qd2, using the same perceptual function, if we are to see the net disturbance at the second level correctly.

Considering only the first-level system, we still have the reference signal determining the input quantity, now qi1b. This means that the output of the second-level system is, as far as second-level control is concerned, not qo1 but qi1b. The input quantity of the first-level system, not the output quantity, will appear to be the action of the second-level system. If qi1a is disturbed, qi1b will change to oppose the effect on the second-level perceptual signal. But it is also true that if qi1b is disturbed, qi1a will change to oppose the disturbance of the second-level input function -- although qi1b, being under control itself, will not give way much to disturbances.

Therefore:

(c) The total disturbance, composed of d1a and d1b, and the second-level reference signal, jointly determine the second-level output, which translates into qi1b in this case.

(d) The second-level reference signal and the second-level input function determine the second-level input signal, which means

$$I2^{-1}(qi2a,qi2b) = r2$$

The appearance will be that an abstract variable composed of qi1a and qi1b will be exclusively determined by the second-level reference signal. At the same time, the output quantity of this system will appear to be qi1b, and it will be jointly determined by the second-level reference signal and both disturbances.

These are fairly complex and subtle relationships. Understanding them requires seeing how control system vary their outputs to maintain their inputs at preselected levels or states, at the same time automatically resisting the effects of disturbances on those inputs.

At any level of interpretation, statements (a) and (b) will hold true -- but with many systems at each level, each level has to be considered anew. When a single control system at one level receives reference signals from several higher systems, there can be no simple relationship between disturbances of a given higher-level perception and the resulting change in the lower-level net reference signal.

This is why I don't think there is much point in trying to express the relationships of control in the familiar language of determination, influence, and causation -- particularly not in terms of causation. In speaking of one simple control system, I have always used causal terms in speaking of illusions: it SEEMS that one variable is causing another to change, but in reality, in a control system, the pathways of causation are quite different from what they seem, and are circular instead of lineal.

I think the closest we can come are the two basic statements:

(a) The reference signal and input function determine the state of the input quantity in the environment.

(b) The reference signal and external disturbances jointly determine the output quantity.

Both of these statements describe apparent causal relationships, which are different from those that actually exist in the control system. That is, the "determination" takes place through a path different from the one that appears to exist. These two statements describe appearances, but not the actual organization of a control system. Both are deductions about how control behavior will appear to a naive observer, based on the assumption that we are observing an ideal control system.

Best, Bill P.

Date: Wed Jun 17, 1992 7:01 am PST
Subject: Nevin Misc.

[From Bill Powers (920617.0800)]

Bruce Nevin (920616) --

>An interesting view of how anarchism works is in Ursula LeGuin's novel
>_The Dispossessed_. A fine writer.

Yes, indeed. This novel is worth study as an illustration of the way system concepts rule one's whole organization. It has long been my favorite LeGuin. I've started a re-read -- thanks.

>I thought that addictive substances grabbed you by the intrinsic error.

What they do (some of them) is give you the same feeling you would get if you sank a 50-foot basket to win the last game of the playoffs. Only without actually accomplishing anything.

Carl Sagan on primates:

> The perks of being an alpha male entail certain obligations. In return for deference and respect, for preferential sexual access to ovulating females and for deluxe dining privileges, he must render services to the community, both practical and symbolic. He adopts an impressive demeanor, even something approaching pomp, in part because his subordinates demand it of him. They crave reassurance. They are natural followers. They have an irresistable need to be led.

This is an example of misapplied categories, or homocentrism. Obligations, deference, respect, services to the community, and the idea of subordinates demanding something of a leader (etc., etc.) are interpretations by a human observer. To use such interpretations is to miss the opportunity to guess what the actual controlled variables and means of control are. You could speak exactly the same way of ants, but then the projection of human qualities would be too obvious. The obverse of this objection is that using such terms for human interactions is probably also to miss an opportunity to see how control is working.

>I have several times posed the question, how can an ECS controlling for
>perception of a word (or category) distinguish between tokens that have
>the same referent and those that do not.

This is more of a problem for understanding speech than for producing it, at least according to my hypotheses about language. Talking with J. Marvin Brown (and eating his wife's cooking), it occurred to me to wonder whether the word for "hot" indicated, in Thai, both temperature and what we call a "hot" taste. The answer is that no, these are different words in Thai. This interested me, because to me, the sensations evoked by high-temperature food seem similar to those evoked by certain peppery spices. I had always taken the English usage of "hot" to indicate this perceptual similarity. Now I see that it's just a failure to make a perceptual discrimination.

This isn't quite what you're talking about, but it's similar. When I'm producing speech, I know in advance the meaning I want to produce, so there's no problem in my own understanding of the words I use. I can say "John is supposed to have said that, but that's just John's opinion." (or "his" opinion). There's no problem in my mind with the two Johns -- I'm still meaning the same nonverbal person, whatever word I use. A problem arises only when I DON'T mean the same John -- when John Brown is supposed to have said the thing, but that is only John Smith's opinion. I could still use the same sentence and know what I mean, but it would be clear to me that a listener might have a problem sorting it out as I mean it.

So I agree:

>One approach that I have not mentioned is to suppose that a word (or a
>category) has the same reference as a repetition of it has unless they
>are explicitly differentiated.

(I distinguish between "referent" -- the perception pointed to by a word -- and "reference" -- the state of a perception that is desired).

The differentiation involves not just using alternative words, but introducing different perceptual attributes, as you say:

>An occurrance of a word may be differentiated from another occurrance >of the same word by asserting something of the first and denying the >same of the second.

The "something" is the different perceptual attribute: you take the red apple and give me the yellow one. Or (in terms of assertion/denial), you take an apple that's red and give me one that's not.

>So previously the question was how to capture the perception of >sameness nonverbally; now it is how to capture the perception of >difference between two tokens of the same perceptual signal.

I don't think this is ever a problem for the speaker. The perceptual meaning comes first; then the attempt to describe it using words. The problem is for the speaker to detect possible wrong meanings that a listener might get (Martin Taylor's "model of the listener"). You can think "He sure didn't like how he was taking her attitude toward her boyfriend" and have no problem keeping the four different people straight. But if you say that to someone else, there's no telling what meaning the other person will construct on that base. The differentiations needed can be accomplished by describing different perceptions that go with the different people: The fat guy sure didn't like the way the thin guy was taking the black girl's attitude toward the grandmother's boyfriend. Same meaning as far as the speaker is concerned; much clearer to the listener.

Sameness and difference don't have to be captured nonverbally. If the same perceptual function responds the same way while other perceptions change, and you're talking about the one that didn't change, you say there's "sameness." If you're talking about the ones that did change, you say there's a "difference." The changes or lack thereof come first; then the attempt to describe them.

>Any insights?

Have you ever known me not to have an opinion? Oh, you said insights.

>I want to work up some ideas on the reduction of logic and the program
>level to simpler terms based on G. Spencer Brown's _Laws of Form_.

See "Control theory, constructivism, and autopoiesis" in Living Control Systems II, p. 179: the section called "Cleaving the continuum." My thesis is that the concept of creating complementary categories by "strokes" belongs at the category level of the HPCT model, but that there is also sequence and logic operating tacitly in the background, and that the "continuum" is created by lower level perceptual systems. At the lowest level, intensities, we begin with the maximum possible differentiation: every intensity signal is an independent entity. You can't cleave the continuum until you've created it.

Nice to hear you vocalizing again. Coming to the meeting? Deadline for registrations is July 13. $220 for everything but beer (includes membership). You've met a few of the CSG: everyone else is exactly the same, only different.

Best, Bill P.

Date: Wed Jun 17, 1992 8:43 am PST
Subject: Machine Learning List

A couple of people asked me about the Machine Learning List.
It turned out my source by that name was a BBN-local distribution
list. However, here is information about another:

-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-
The Machine Learning List is moderated. Contributions should be relevant to the scientific study of machine learning. Mail contributions to ml@ics.uci.edu.
Mail requests to be added or deleted to ml-request@ics.uci.edu. Back issues may be FTP'd from ics.uci.edu in pub/ml-list/V<X>/<N> or N.Z where X and N are the volume and number of the issue; ID: anonymous PASSWORD: <your mail address>

Bruce bn@bbn.com

Date: Wed Jun 17, 1992 11:07 am PST
Subject: more to Andy P and Tom B

[From: Bruce Nevin (Wed 920417 13:28:02)]

(Andy Papanicolaou and Tom Bourbon [920604 13:50]) --

Having during this morning's commute at last worked my way back as far as the post to which I responded second-hand, by way of Bill's response to it, I see that you had anticipated much of what I had to say. My apology if I seemed obtuse is that I was (unnecessarily) ignorant.

Quoting (with some reformatting to help spatial organizers like me), you want

>to construct a CT model that would account for . . .
> acquiring [new speech sounds]
> or learning the skill [of pronouncing new speech sounds]
> or forming the habit of pronouncing new speech sounds
> or "reorganizing" to that effect . . .

First off, in the Japanese r/l case what you want is not a new (phonetic) speech sound but a new (phonemic) distinction, a new contrast. There is a range of speech sounds (better: a pair of ranges) that speakers of English find admissible as contrasting r with l. So long as these are consistently distinguishable through respective ranges of allophonic variation, and so long as each is distinguished from all other "phonemes" (reason for quotes presently), it works as a dialect of English.

If a speaker shifts r/l along one or more parameters of (phonetic) control of speech sounds, then the speaker must also shift other contrasts that have neighboring values along those parameters in order to preserve their contrasts with r and/or l.

I am getting in a bit of a muddle here, and this is the reason (also the reason for putting "phonemes" in quotes): it is the relationship of contrast that is primary, not the relata, the phonetic speech sounds instantiating the contrast in particular utterances. But it is difficult for us to talk about a relationship other than in terms of its relata. English just doesn't work that way.

Before we run off into a metaphysical discussion of how the relata must logically be prior etc., reflect on what is being contrasted. People control for keeping words distinct from one another. (More generally, morphemes, including thereby argument indicators, operator indicators, and reduced forms of words, but that does not matter for this discussion.) My hunch is that in the process of learning language control for maintaining the contrast of differently pronounced syllables and semisyllables and perhaps segments and features is arrived at analytically as means for controlling the contrast of words. At the level of words, the relata are prior. However, the contrast of words is complex in ways that invites analysis leading to control of syllables, semisyllables, and perhaps segments and features.

I think acoustic feedback is many times more important than tactile or kinesthetic for this control. This fits with the social status of language, as Roman Jakobson pointed out long ago. However many have worked long and hard on articulatory gestures as phonological primes. the most important in my opinion is that at Haskins Laboratory. See for example and references:

Browman, Catherine P. and Louis Goldstein. 1989. Articulatory gestures as phonological units. _Phonology_ 6:201-251.

The experiment that I would like to run, but I do not have the equipment or the programming expertise, would alter the auditory feedback of the participant's own speech in real time. Suppose we distorted the frequency range of the second formant, gradually bringing [U] of "foot" closer to the the [u] of "food." (The degree of distortion would ideally taper off at the boundaries of the envelope of distortion, to avoid discontinuities and to make transients with adjacent consonants sound more natural.) I predict that the speaker would gradually shift the F2 of the [U] vowel lower and lower, closer to [o] and [^] and that the speaker would not be aware of doing so. This would make a difference in pronunciation quite noticeable to one listening to the unaltered audio.

This would not be easy. It is even more difficult to distort the audio feedback for consonants. There might be a way to affect VOT (voice onset time), that is, the delay between laryngeal and oral gestures. This makes the difference between /b/ and /p/ in English, and the difference between, say, English b/p and spanish b/p.

This then is well into the subphonemic aspect of your question, exemplified by learning a native or foreign "accent," where the term "speech sounds" is more appropriate. But on that level, one is controlling other things than contrast. This is not learning the language per se, but rather the gestural use of the language for things like affect, self image, and group identification. An example is Joel's Korean student of English who could speak flawless Surfer at the expense of feeling unacceptably silly and inane to himself, or the man who would not learn proper French pronunciation because he could not bring himself to hold his mouth in that prissy way.

I'm not meaning to dance all around your question, I'm trying to suggest different ways of putting the question--and that it is considerably more complex than has been thought.

Bruce bn@bbn.com

Date: Wed Jun 17, 1992 11:19 am PST
Subject: Re: car event

[From: Bruce Nevin (Wed 920417 14:44:47)]

(David Goldstein (920604)) --

>Was this just a coincidence? Psychic phenomena? Was I in imagination
>mode and this played a part in the incident? What are your
>speculations?

The standard "occult" explanation (in PCT-ish terms) is that control of one's perceptual universe has a creative relation to physical reality. This is why for example the Indian people I know won't wear seatbelts, since to do so out of worrying about accidents would contribute to causing an accident. (My reply to them now would be that I do it automatically, without the worry, and that I started doing it as an example so my children would see it as an automatic part of riding in a car, because I didn't want bouncing around the car if I cornered or braked hard.)

Within that mode of explanation, there are constraints on our creative physical effects, based in interpsychic agreements about the way our shared physical world works.

This mode of explanation accounts for the experience of adherents of a "scientific" (i.e. current engineering) explanation. Such adherants do not experience exceptional "occult" or "psychic" events because their use of the creative capacity of their control of their perceptions does not countenance them.

Disproving this mode of explanation is not easy. We know expectation can color or even engender perception. We only presume that that-which-is-perceived has a stability that remains autonomous from our perceptions of it. In other words, can we agree about Boss Reality?

Bruce bn@bbn.com

Date: Wed Jun 17, 1992 12:27 pm PST
Subject: Re: more to Andy P and Tom B

[Martin Taylor 920617 16:00]
(Bruce Nevin Wed 920417 13:28:02) 920617?

>First off, in the Japanese r/l case what you want is not a new
>(phonetic) speech sound but a new (phonemic) distinction, a new
>contrast. There is a range of speech sounds (better: a pair of ranges)
>that speakers of English find admissible as contrasting r with l. So
>long as these are consistently distinguishable through respective ranges
>of allophonic variation, and so long as each is distinguished from all
>other "phonemes" (reason for quotes presently), it works as a dialect of

>English.

The distinction between phonemic and phonetic is very important, as Bruce says.
Koreans do distinguish r and l phonetically, but not phonemically. What we hear as r is the version they tend to use at the beginning of syllables, and what we hear as l they tend to use at the end of syllables (or maybe it's the reverse). There is no phonemic contrast, but if the wrong one is used, it sounds wrong and a Korean would not make the mistake. But they have a lot of difficulty perceiving the difference in English, where the contrast is phonemic. Even after 35 years in an English speaking environment, my wife still mixes them up occasionally, even in writing (she writes testbooks on
language, in English). Sometimes I have to make quite an exaggerated distinction in pronouncing two words to her, if they differ only by the r-l difference. On the other hand, I cannot tell whether she says Bulgogi or Prugogi. I hear the latter, but she says she speaks the former (it means a kind of braised beef). Notice that this problem includes a question of where in the syllable the r/l comes, as well as which one it is.

>I am getting in a bit of a muddle here, and this is the reason (also the
>reason for putting "phonemes" in quotes): it is the relationship of
>contrast that is primary, not the relata, the phonetic speech sounds
>instantiating the contrast in particular utterances. But it is
>difficult for us to talk about a relationship other than in terms of its
>relata. English just doesn't work that way.

Yes, I think this is true at all levels of language. It is the situation-specific relationships that require discrimination that matter. We ordinarily do not repeat a discriminating identifier if we want to refer to the same thing--we use anaphora to say "the same as before". To repeat the discriminator
is to say "not the same as before." Which is why I had a problem with Bruce's example, which I don't have to hand for a direct quote, but it went something like: "John says he thinks that, but that's John's opinion." It takes the pragmatic context to override the initial signal that the second John is different from the first. Indeed, on consideration, I think that the second "John" IS connotatively different, though denotatively the same as the first. The first John, in a neutral way, selects the person from many who could have been identified. The second seems to emphasize the discrimination, and to say that ONLY John could have such an opinion, all right-thinking people having
the opposite opinion.

At a much lower level, the second and subsequent uses of a content word in a discourse are much less intelligibly spoken than the first. (This is tested by cutting the occurrences out of a recording and using them in intelligibility
tests). I suspect that if the word were to be spoken as clearly on the second occurrence, the clarity would be a signal that there is a new intent in using the word, not the same as before. Reduced forms at all levels tend to indicate
that the talker models the listener as being able to provide the required information from memory (or other sources), whereas full forms indicate that

the talker believes that the listener does not have the information, and therefore that this presentation provides something new.

Perhaps this goes a little way to answering Bruce's query about the detection of sameness, at least in conversation?

Martin

Date: Wed Jun 17, 1992 2:39 pm PST
Subject: Phoneme relationships

[From Bill Powers (920617.1600)] Bruce Nevin(920617) --

A thought: if relata are "logically prior" to relationships, HPCT would say they are of a lower level (i.e., you can't have relationships without relata, perceptions of lower order, but the reverse is not true). But CONTROL of relationships would still take precedence over production of specific relata, because to CONTROL a higher-level variable, you must VARY lower-level ones. Thus maintaining the distance between phonemes, a relationship, might well entail altering the individual phonemes, which are configurations or transitions.

Best, Bill P.
Date: Thu Jun 18, 1992 5:53 am PST
Subject: same cat

[From: Bruce Nevin (Thu 920418 08:03:40)]

Bill (920612.1300) --

The discussion of "same" perception under symmetric reflections bears in a perhaps interesting way on my metalanguage concerns.

>If a perceptual function sees something as "the same" when it's turned
>over, I don't think it reports it as "the same." It simply goes on
>reporting it without any change. OTHER perceptual functions might see
>changes: apparent size, velocity, . . .

Underlying the notion "same" is the notion of an individual, perduring through time. The timeline of an individual comprises a history in memory (plus imagination) and a future projected in imagination. This timeline extends on either temporal "side" of the present. The present of course comprises real-time perceptual input plus imagination.

Some perceptions remain the same along this timeline, for example, the category perceptions <pet> and <cat>. (I am using <, > for nonverbal perceptions and reserving " for words.) What does this mean? Well, suppose there is just one ECS that functions as a cat recognizer. Signals from this one ECS provide input to a number of other ECSs. Some of them are in the present ("Miaow!", staring eyes, rubbing against chair). Some are in memory (childhood memory of adult saying "Yes, it's time to feed you, isn't it?" to a cat, memories of oneself saying similar things, memories of cats stopping this assertion of dependency and running to eat). Some are in imagination (she'll go on doing this until I stop writing and go feed her). All of this could be the basis for discourse about the present situation. In such discourse, most repetitions of the word "cat" having the same referent can be reduced to things like pronouns or to zero.

We can suppose that the business of providing input to another ECS is equivalent (graphically) to a line in a graph, a mesh or net, where the vertices or nodes are words associated with each such ECS. (We have some rough suggestions of what "associated with" might mean.) Different ECS nodes may change the level or manner of their participation in such a mesh over time, where the three levels or manners I know about are real time, memory, and the imagination loop. We believe we can tell the difference--for example, when something expected actually occurs, or when an occurrence turns out to be familiar (the "same" association of "miaou" perception with <cat> perception).

In this way of representing things, how is it that the ECS for <cat> or the word "cat" can have two or more different referents? There must be several distinct meshes of associated perceptions. These may intersect almost entirely (after all, most cats are alike in most respects). We then pay attention most to those attributes (associated perceptions) that differentiate them. The ECS for <cat> and the ECS for "cat" and that for <tail> and <miaow> and a host of other perceptions just go on reporting their perceptions without change. But perhaps some of them differ as to the level or manner of their input. For example, say the ECS for <tail> is in real time for one cat and in imagination for a second, because his tail is presently out of your sight, in memory for a third who lost her tail in an accident, and inactive (as a distinguishing attribute) for cat #4, a manx. Can we maintain this degree of discrimination?

An alternative is an indefinite number of ECSs for <tail>, etc.

Now what about a category like <cat> or <tail>? The basis for categorization is analogy. A category perception looks like the outcome of an analytical process abstracting attributes common to exemplars. What if we start with one or more exemplars, where an exemplar is an associative network or mesh as sketched above? An analogical process would check current perceptions for fit with the mesh established for familiar exemplars. A remembered attribute becomes the basis for imagination. When in doubt, we explicitly test imagined attributes, especially those attributes that distinguish one category from another (or one remembered individual from others). When not in doubt, we implicitly test some imagined attributes (though not necessarily those that are crucial for distinctions) simply by projecting a future timeline for the present individual and acting on those predictions.

Language comes in as a set of associative hooks for categories established by previous generations. The categories and the verbal hooks for them are explicitly taught to children, and children are eager to learn them because the categories facilitate control and more importantly because the words help the children to elicit the cooperation of others in accomplishing their aims.

The nonlinear mesh sketched here provides a base for linearizing alternative discourses about the subject matter of the mesh. In my 1969 MA thesis I called this periphrasis, as distinct from paraphrase within the sentence. I think this view is quite congenial with yours, Penni.

Gotta quit for now. Bruce bn@bbn.com

Date: Thu Jun 18, 1992 6:54 am PST
Subject: Re: same cat

[Martin Taylor 920618 10:00]
(Bruce Nevin Thu 920418 08:03:40)

Bruce brings up a problem that has long bedevilled neural network students: If the net can recognize an item of class X, how can it recognize that there are two items of class X and keep track of them. I know of no satisfactory solution, though there have been many proposals. Perhaps someone more up to date on the neural net literature can provide references. (This is relevant, because on the input side, the interconnections of the perceptual functions of the ECSs is exactly a feed-forward neural network. If the perceptual functions are limited to weighted summations followed by nonlinearity, the input connections form a multi-layer perceptron.)

In tracking identically shaped objects moving randomly in a visual field cluttered with other objects of the same kind, humans seem able to track three objects perfectly, four with difficulty, if the objects are moving too fast to allow the observer to shift attention from one to another. I would guess that similar limitations occur at higher levels of abstraction. (I may be out by one on "three" and "four", since I am remembering a talk by Zenon Pylyshyn this February).

>An alternative is an indefinite number of ECSs for <tail>, etc.

In our old study on reversing figures, we concluded that one of our observers had 26, and the other 33, units devoted to the perception of the figure orientation. Those units could be ECSs.

None of this solves Bruce's problem, but some linkage to artificial and to human systems might be helpful in situating the problem more securely.

Martin

Date: Thu Jun 18, 1992 8:50 am PST
Subject: same cat

[From Bill Powers (920618.0800)] Bruce Nevin (920618) --

>Underlying the notion "same" is the notion of an individual, perduring
>through time. The timeline of an individual comprises a history in
>memory (plus imagination) and a future projected in imagination. This

>timeline extends on either temporal "side" of the present. The present
>of course comprises real-time perceptual input plus imagination.

This seems to bring up a lot of levels of perception, from categories through system concepts (the concept we refer to as "an individual"). I'm not sure what to do with "timeline." This would seem to entail the sequencing of memories and imagined extensions of the present. Perception of sequence or ordering in terms of a relation between memory and current perception would create time, wouldn't it? This would be nearly the same as the perception of causation.

>In this way of representing things, how is it that the ECS for <cat> or
>the word "cat" can have two or more different referents? There must be
>several distinct meshes of associated perceptions. These may intersect
>almost entirely (after all, most cats are alike in most respects).

Let me try something:

```
                    <cat> <-----------------> "cat"
                      ^
                      |
            [Category perception]
            ^      ^      ^     ^
            |      |      |     |
        <cat1>   <cat2> <cat3> <cat4> <---> "names of cats"
```

If we consider only words, then it seems that "cat" somehow refers to "Ginger," "Aleptic," "Piewacket", and so on, the names of specific cats perceived as discriminable configurations. In terms of nonverbal perceptions, however, we must first be able to perceive different cats as different, so we have a set of perceptions that are NOT alike, one for each cat that we can distinguish from other cats. Then we create a category perception that responds to all those different cats with a single perceptual signal, <cat>, which is called "cat" or "cats". As you can see, I'm making an attempt to treat words as labels for nonverbal perceptions, with the actual relationships existing among the nonverbal perceptions and explaining the apparent relationships between the words.

With this picture there's no problem with having the name of a category-perception "refer" to different perceptions of lower order. Controlling for the perception <cat> entails finding any action that will bring into perception (at a lower level) any one or more of <cat1, cat2, ... catn>. To fulfil the verbal request to show me a "cat", you first translate "cat" into <cat>, then take whatever action will find one of the lower-level perceptions that is in this category, matching the reference signal <cat>. If I agree that the item to which you're pointing is a <cat>, I will probably also call it "cat" (although I might say "feline"), and we will reach verbal agreement: You have fulfilled my request to "show me a cat". Even if I call this category "gato", we will agree -- and in fact I will learn that what I mean by "gato" is what you mean by "cat." So in my diagram above, the items on the left are meanings, and the items on the right are words that refer to the meanings.

This is, of course, a very simple case. I think, though, that even if we end up with your concept of two "meshes", the best way to start is by finding the simplest cases we can and gradually adding complexity as we find the need for it. If we start with this simple naming relationship and see how much can be accomplished with it alone, a certain range of linguistic phenomena will fall into place. Then it might turn out that we could treat "words" as <words> in a similar framework, and come up with the same level of explanation of metalanguages. When this has gone as far as it can, it's time to start looking at the relationship between <structure> and "structure", and so on.

I see what you're doing, but I'm uncomfortable with it because it draws on so much that is outside the HPCT framework without questioning it -- concepts like "timeline" and "individual," for example. If we're going to try to build an HPCT model of language, it seems to me that we should try to formalize all the important informal concepts, or concepts based on other approaches, that are used in the construction. I'm not a mathematician, but it seems to me that the concepts of HPCT have to form a "group", in that the legitimate operations on the entities of the theory ought to leave us still within the theory.

There's a mode of theorizing that I call "truthsaying." What you try to do, at least as long as you can keep it up, which may be only five minutes, is to make a series of statements about the subject matter that are ALWAYS ABSOLUTELY TRUE as far as you can tell. This means leaving out everything that's just a possibility or a proposal or a generalization, or that's true only some of the time or of some people, or that might in some conceivable (but reasonable) way be false. What you end up saying, of course, is trivial and obvious. "People say words." "The meaning of a word is something I can perceive." "A word is a sound or a visual configuration." That sort of thing. Really dumb, but really true. When you have collected a lot of such banal observations, you then try to say something equally true about them. You just trust that if you keep trying to truthsay, something will pop out that is obviously true that you haven't thought of before.

This has worked for me when all else has failed. Usually the result has turned out to be not much, and not even necessarily true on later reflection or investigation, but it has invariably left me going in some new and useful direction; usually the problem ends up solved. Maybe it would have ended up solved, anyway. But truthsaying forces you to keep it simple and short, and there's a certain hypnotic satisfaction in the process once you get it started.

Language is a very complex subject. But there must be things we can say about it that are always, without exception and without doubt or controversy, true. That would seem like a good place to start.

Best, Bill P.

Date: Thu Jun 18, 1992 10:22 am PST
Subject: projection; Lieberman

[From: Bruce Nevin (Thu 920418 13:57:59)]

There was a query in the linguist digest a while ago about voice quality. She was talking with a voice teacher about why college instructors are such poor speakers.

>He said . . . 'The main problem is that they speak down here [indicating the >ventral pharyngeal wall] rather than in the front of their mouths'. . . . Can >anyone interpret this expression? More generally, is there any source which >interprets in phonetic terms the weird but apparently effective instructions of >speech and singing teachers, such as 'sing with your forehead'?

All of this refers to a perceived focus of resonance. I think Liberman and Blumstein talk about it a bit. A "covered" tone or "chest" tone seems to have higher frequencies damped, and so is less "penetrating". A "covered" tone seems to me to be produced by expansion of the pharynx ("yawning"). A "head" tone seems to be a resonance through the sinus cavities that may be helped or initiated by something like pharyngealization. (It feels something like nasalization, but visual inspection in a mirror suggests only contraction of the pharynx is involved.) Differences of vowel quality may be involved.

This fits with my understanding that acoustic feedback is more important than kinesthetic or tactile feedback for control of language. This is not language per se, of course, but it is closely related.

People who can control these distinctions can't tell you how they do it (the articulatory actions), they can only describe the result. Sound familiar? And the result is described in subjective terms: how one's voice sounds from inside one's own head, which is different of course from how the teacher's voice sounds to the student in a demonstration.

Bill, I think your suggestion of providing alternative means for controlling an audible distinction, as part of the first step of learning the distinction, is an excellent and important one. It might apply to a wide range of learning situations where characteristics of automatized action (timing, sequence, etc) already established for different but analogous categories might foreclose the learning process--in effect imposing the established but now inappropriate categorization by the back door of its automatized behavioral outputs!

Just in passing, Bill, I think you will be intrigued by what Lieberman has to say in _The Biology and Evolution of Language_, particularly in his chapter 3, Automatization and Syntax. This has been on my stack for some time to get back to. He suggests in that chapter for example that structures for motor control were adapted for syntax. It is interesting that lesions in Wernicke's area result in problems with speech perception and semantics (various types of Wernicke's aphasia) and those in Broca's area result in problems with articulation and with the issues of "grammar" (i.e. reductions) and interruptions (as for relative clause) that Genie could not manage. Wernicke's area is an auditory association region connected with the auditory system (Heschel's gyrus); Broca's area is an "association" region similarly adjacent to the motor control area of the brain (ibid. 25-29, 47). He has some interesting discussion of hydrocephaly evidence against received wisdom that the cortex is the essential and primary neural seat of intelligence, e.g. discussion of the following quotation:

There's a young student at [Sheffield University, England] who has an IQ of 126, has gained a first-class honors degree in mathematics, and is socially completely normal. And yet the boy has virtually no brain. We did a brain scan on him; we saw that instead of the normal 4.5 centimeter thickness of brain tissue between the ventricles and the cortical surface, there was just a thin layer of mantle measuring a millimeter or so. His cranium is filled mainly with crebrospinal fluid. (J. Lorber, Is your brain really necessary? Research news. _Science_ 210:1232-1234, quote from p. 1232)

Bruce bn@bbn.com

Date: Thu Jun 18, 1992 11:46 am PST
Subject: intro

Hello-

I'm not on the list, I get CSG via the Usenet gateway.

Very interesting topic. I have several comments which require some research, but here's a naive one:

You folks complain about resistance to the control paradigm. I have one suggestion: throughout Western intellectual history we find the pattern of assuming that the human brain works via the same mechanism as the most complex gizmo that we know how to build. The most complex and interesting engineering that the Romans knew was water distribution, and you'll find musings among Roman "natural philosophers" that the brain has blood and tubing so therefore it must be a marvelously complex hydraulic system. In the 17th & 18th centuries it was clocks, the 19th engines and then wow! the telephone exchange.

Control systems are not the latest spiffy gizmo, therefore the brain can't possibly work like one. It has to be something more complex, like digital computers or these new neural network thingies. Or maybe Artifical Life. Yeah! It has to be a new mechanism whose mysteries we haven't plumbed, not those old boring control systems.

The AI biz has had this long history of inventing new software technologies which, when applied, are suddenly not AI but just "clever programming". The brain can't possibly function like some technology which has lost its mystery.

Lance Norskog

Date: Thu Jun 18, 1992 3:13 pm PST
Subject: Multiple X perception; brain deficits; CT is old hat

[From Bill Powers (920618.1400)]

Martin Taylor (920618) --

>If the net can recognize an item of class X, how can it recognize that >there are two items of class X and keep track of them. ... In tracking >identically shaped objects moving randomly in a visual field cluttered >with other objects of the same kind, humans seem able to track three >objects perfectly, four with difficulty ..

So this gives us a starting point: (roughly) four objects is the limit, showing that whatever the system, it isn't some general principle applicable to n objects (like a hologram). There's a strategy and it has to time-share among the tracked objects.

Let's try changing the question: not "how can it recognize that there are two (n = 1-4) items of class X" but "how can it recognize that there are two items?" and "how can it recognize that a given item is of class X?" The first questions concerns counting, the second concerns perception of class membership.

Counting is probably a phenomenon involving the sequence level and up. So one strategy is: set counter to zero; if the item yields a perception of class X, count it; go to next item. In place of counting (if you want a parallel-processing system), you could say that the magnitude of the X signal depends on how many items are present that fit that category. If we can reliably detect only 1 to 4 such items, this says that a single item yields about 25% of the maximum signal from that category detector. From the magnitude we estimate the number.

The main thing is to separate perception of number from perception of class membership, then let a higher system put them together. When we SPEAK of such situations, we collapse many perceptions at many levels into one package. Just think of all the levels of perception involved in "Will the last person out please turn off the lights?"

Bruce Nevin (920618) --

>It is interesting that lesions in Wernicke's area result in problems >with speech perception and semantics (various types of Wernicke's >aphasia) and those in Broca's area result in problems with articulation >and with the issues of "grammar" (i.e. reductions) and interruptions >(as for relative clause) that Genie could not manage.

I've seen such reports. But the tests that are used really can't tell the difference between mishandling speech per se and mishandling meanings. This is why I'd like to see the aphasias studied with parallel control-system tests. It isn't necessary to talk about a relationship like "above" or "beside" to control it -- make it match a reference relationship, for example. So you can find out if something is being perceived as a nonverbal experience independently of whether a person is able to describe that kind of experience. Many aphasics can control for variables that they can't talk about.

I've also run across references to the "no-brainer" case that you cite:

There's a young student at [Sheffield University, England] who has an IQ of 126, has gained a first-class honors degree in mathematics,

and is socially completely normal. And yet the boy has virtually no
brain. We did a brain scan on him; we saw that instead of the normal 4.5 centimeter thickness of brain tissue between the ventricles and the cortical surface, there was just a thin layer of mantle measuring a millimeter or so.

This says something informative about either brains or first-class honors
degrees in mathematics. As to being socially completely normal, I'm not sure it takes a full complement of brains to be that.

More seriously, this is a strong argument for extreme plasticity in the cortex -- for reorganization as a major factor in development. I would have to guess that this condition was present from birth -- it would be hard to imagine a brain developing normally and THEN losing so much of its bulk, while continuing to function in a way that seems normal. If my guess is right, then we would expect whatever control systems did develop in that millimeter or so would have reorganized to control variables at all the necessary levels -- even if not as many variables as otherwise.

Another thought: it is possible for higher systems to get around deficits at intermediate levels (as some data about aphasia show). Deafferented monkeys can still reach out and touch a visual target if they can see where their hands are (after a sufficient period of practice after surgery). They never can do this, especially in the presence of disturbances, as well as if they were intact, but they can regain some control. Unfortunately, the people who do such studies never do real control-system tests, so there's no way to know just how much control is actually regained.

Incidentally, a lot of deafferentation data is made suspect by the period of recovery after surgery. During this time, animals can easily learn to use other feedback paths: sensations from hairs on the skin, or skin- pressure receptors, or shifts in balance, and so on. Such things could account entirely for the clumsy "control" that is seen in such studies after deafferentation. They're not looking at the same nervous system after the surgery.

It's interesting that this student chose mathematics. If the primary subject had not been logical symbol-manipulation but, say, taxonomy or horticulture, could he have done as well?

Lance Norskog (920618) --

Thanks for speaking up!

>Control systems are not the latest spiffy gizmo, therefore the brain
>can't possibly work like one.

You've got it in one try. About 30 years ago, the head honchos of behavioral science, who never understood control theory anyway, decided that "servomechanism" models were out of date, and went on to bigger and better things (one person who gave up on control theory was the then president of the American Society for Cybernetics). This has been the story of behavioral science: a new fad every decade or less, with all the work done on previous fads being forgotten. We spent some time on this subject a year or so ago on CSGnet -- "trendy science." Every time someone comes up with a new mathematical spiffy gizmo, every one seizes on it as the answer to their prayers and starts applying it to every problem that's failed to be solved (most of them, as nobody really sticks with any approach long enough to see if it will really work). Enthusiasm, earnest effort, disillusionment, discouragement, oblivion.

What do you think of this? "Control theory contains the first new conception of behavior since Descartes."

Tell us something about yourself -- you might as well subscribe to this list, while you're at it. I can tell you are a Right Thinker.

Best to all, Bill P.

Date: Thu Jun 18, 1992 3:25 pm PST
Subject: Flood report

From Greg Williams (920618)

This morning Black Lick Creek received, in about one hour, approximately five inches of rain. Our old house was slightly flooded, our new house is drenched but basically OK, our barn was about half flooded, and the yard needs a LOT of work. But we (including pets) are OK.

I'm going to be inordinately busy for awhile. Bill, I hope you will accept a slight setback on editing the arm paper. Gary, I'd greatly appreciate your disconnecting me from the Net temporarily, until I tell you to put me back on. I can get the log files from you on what I missed.

I'm still planning to get CLOSED LOOP out on time and to attend the meeting. If you want to reach me, please telephone at 606-332-7606. I won't have much time to check on the net.

Thanks for understanding, everyone.

Greg

P.S. Tom, thanks for going the extra mile on the deafferent references. Will be of great use when I can get back to helping Bill.

Date: Thu Jun 18, 1992 3:27 pm PST
Subject: Re: Multiple X perception; brain deficits; CT is old hat

[Martin Taylor 920618 19:10]
(Bill Powers 920618.1400)

On types of aphasia--the researchers are not so blind as you seem to think:

>Bruce Nevin (920618) --

It is interesting that lesions in Wernicke's area result in problems with speech perception and semantics (various types of Wernicke's aphasia) and those in Broca's area result in problems with articulation and with the issues of "grammar" (i.e. reductions) and interruptions (as for relative clause) that Genie could not manage. I've seen such reports. But the tests that are used really can't tell the difference between mishandling speech per se and mishandling meanings. This is why I'd like to see the aphasias studied with parallel control-system tests. It isn't necessary to talk about a relationship like "above" or "beside" to control it -- make it match a reference relationship, for example.

There are lots of different kinds of tests, including picture matching to word, picture to picture, acting out, labelling action, describing (verbally) pictures or other words ... Many different deficits of abilities have been extracted, and the non-necessity of one for the completion of another can be determined by finding "double dissociation." In one patient, function A is absent but B is present, whereas in another B is absent and A present. One patient may write fluently and correctly, but be unable to read what was written, a few minutes later. Aphasics don't necessarily have trouble reading, but there is a tendency that way. A Broca's aphasic will not be able to use the syntax in text, whereas a Wernicke's aphasic may not make sense of the text. The distinctions are by no means as clear as I make out; each patient is different, and many are so ill that it is hard to test exactly what they can and cannot do. But it is quite unfair to say that the tests used can't tell between mishandling speech and mishandling meanings. That's one thing that they do do, and in addition they tell between mishandling meanings (in the external world) and mishandling functions.

Martin Taylor

Date: Fri Jun 19, 1992 4:19 am PST
Subject: Computer processing of speech

Our organization, the Laboratory of Comparative Human Cognition (LCHC) at the University of California, San Diego, is carrying on a collaborative project with the independent VEGA International Laboratory in Moscow to expand telecommuni- cation contacts between humanities scholars in Russia and their counterparts in the West. As part of this project, VEGA in September 1991 opened a public-access electronic mail address (PSY-PUB) on their premises, through which any humanities scholar in the Russian Academy of Sciences may send and receive messages.

LCHC routinely receives messages from PSY-PUB and forwards them to potential Western partners (both individual scholars and scholarly organizations). We received the following message from Professor

Ivanov of the Linguistics Institute, Russian Academy of Sciences, and are forwarding it to your discussion group in the hope that it may be of interest.

We ask that anyone who responds to Dr. Ivanov please also send a cc to the respective directors of LCHC and Vega, Dr. Michael Cole and Dr. Alexandra Belyaeva. The addresses for Dr.Ivanov, and Drs. Cole and Belyaeva are as follows: Their addresses are as follows:

psy-pub@comlab.vega.msk.su (in the subject line, type "For V.B. Ivanov"; it will be forwarded to him) mcole@weber.ucsd.edu abelyaeva@home.vega.msk.su

Sincerely, Doug Williams

(Message follows)

I, Ivanov Vladimir Borisovich, an instructor in the Laboratory of Experimental Phonetics of the Linguistics Institute of the Russian Academy of Science, wanted to take this opportunity to send a report to our colleagues in connection with the organization of a conference on computer linguistics. I would be very grateful if in the future I would have the possibility of sending electronic mail in connection with research on the computur-processing of speech, discerning multi-language texts, and a multi-language data base.

Thanks ahead of time,
V.B. Ivanov

Date: Fri Jun 19, 1992 11:05 am PST
Subject: chimps

[From: Bruce Nevin (Fri 920419 13:48:06)]

I said (920616):

>One of the problems of individual and community is that mammals and
>probably other animals take their lead from their fellows in setting
>some reference signals. Much of how we do things with words depends
on
>this. There is an interesting survey of primate research in a new
book
>by Carl Sagan and his wife Ann Drayan, _Shadows of Forgotten
Ancestors_.
> [quote from Sagan/Drayan book omitted here]
>A social and political structure that is easily recognizable as
>human-like, complete with palace intrigues and Marshall Dillon style
>enforcement of lawn order, has evolved and is maintained among
>chimpanzees without language as we know it. This is not to deny that
>we do the same with language, and more elaborately by virtue of having
>language. The point for me is that we use language for such purposes
>in parallel with the essential function of language, which is
error-free

>transmission of information. Our use of language for social status,
>political process, communication of affect, and so on is essentially
>gestural, just like our use of facial expression, posture, and so on.

Bill (920617.0800) ignored the issues suggested here because he judged the quotation from the book to be

>an example of misapplied categories, or homocentrism. Obligations,
>deference, respect, services to the community, and the idea of
>subordinates demanding something of a leader (etc., etc.) are
>interpretations by a human observer. To use such interpretations is to
>miss the opportunity to guess what the actual controlled variables and
>means of control are. You could speak exactly the same way of ants, but >then the projection of human qualities would be too obvious. The obverse >of this objection is that using such terms for human interactions is >probably also to miss an opportunity to see how control is working.

I suggest that the projection of our own imagined experience (with its imagined behavioral outputs) onto the observed behavioral outputs of others is an essential preliminary to scientific work. Appropriate scientific method can assure us that hunches and proposals, however we arrive at them, bear a valid relation to direct perceptions. It should not be used to stifle informed speculation, without which we get few hunches and proposals to test.

Imagine the following scene:

The Chief is sitting bolt upright, jaw set, staring confidently into middle distance. The regalia on his head, shoulders and back gives him an even more imposing aspect. Before him crouches a subordinate, in a bow so deep that his gaze must be fixed on the few tufts of grass directly before him. He may even kiss the Chief's feet. Calm and assured, the Chief does not scowl at his nearly prostrate subordinate. Instead, he reaches out and touches him on the shoulder or head. His subordinate slowly rises, reassured. The Chief walks on, touching, patting, hugging, occasionally kissing those he encounters. Many reach out their arms and beg for contact, however brief. Almost all--from highest rank to lowest--are visibly buoyed by the Chief's touch. Anxiety is relieved, perhaps even minor illnesses cured, by the laying on of hands.

The players in this scene are human beings, perhaps in a jungle village, perhaps in a medieval kingdom, perhaps in a meeting of a Mafia "family," perhaps even closer to home. Observing them, we imagine what the experience would be like, for us to be acting like one player or another in this scene. By the uniformitarian hypothesis that underwrites anthropology and all the social sciences and indeed our most mundane essay at everyday communication, we assume assumes") that the perceptions for which we find ourselves controlling in imagination, were we interacting with others as they are interacting together, are indeed the same sorts of perceptions for which they are controlling as we observe them. This is a principal means of determining how we might apply The Test for Controlled Perception in an experimental situation.

The quote I provided was out of context, which was fair neither to Sagan & Drayan or to you, Bill. Here it is again, with somewhat more context provided.

> The alpha male is sitting bolt upright, jaw set, staring confidently into middle distance. The hair on his head, shoulders and back is standing on end, which gives him an even more imposing aspect. Before him crouches a subordinate, in a bow so deep that his gaze must be fixed on the few tufts of grass directly before him.
>
> If these were humans, his posture would be recognized as much more than deference. This is abject submission. This is abasement. This is groveling. The alpha's feet may, in fact, be kissed. The supplicant could be a vanquished provincial chieftain at the foot of the Chinese or Ottoman emperor, or a 10th-century Catholic priest before the Bishop of Rome, or an awed ambassador of a tributary people in the presence of Pharaoh.
>
> Calm and assured, the alpha male does not scowl at his nearly prostrate subordinate. Instead, he reaches out and touches him on the shoulder or head. The lower-ranking male slowly rises, reassured. Alpha ambles off, touching, patting, hugging, occasionally kissing those he encounters. Many reach out their amrs and beg for contact, however brief. Almost all--from highest rank to lowest--are visibly buoyed by the King's touch. Anxiety is relieved, perhaps even minor illnesses cured, by the laying on of hands.
>
> Regal touching, one after the other in a sea of outstretched hands, seems familiar enough to us--reminiscent of, say, the President striding down the central aisle of the House of Representatives just before the State of the Union address, especially when he's riding high in the polls.
>
> <Omitted: Discussion of close genetic relation of chimpanzees and humans, and of some apparently homologous aspects of chimpanzee and human social life.>
>
> The chimpanzee alpha male will intervene to prevent conflict--especially between hotheaded young males, pumped up on testosterone, or when aggression is directed at infants or juveniles. Sometimes a withering glance will suffice. Sometimes the alpha will charge the

pair and force them apart. Generally he approaches with a swagger, arms akimbo. It's hard not to see here the rudiments of gonvernment administration of justice.

The perks of being an alpha male entail certain obligations. In return fo deference and respect, for preferential sexual access to ovulating females and for deluxe dining privileges, he must render services to the community, both practical and symbolic. He adopts an impressive demeanor, even something approaching pomp, in part because his subordinates demand it of him. They crave reassurance. They are natural followers. They have an irresistable need to be led.

The anger of a high-ranking male is fearsome. He may charge, intimidate, and tear branches from trees. He exaggerates his size and fierceness and displays the weapons that he will bring to bear if the adversary does not submit. These displays are used for keeping more junior males in line. Displays may serve as a response to a challenge, or just as a general reminder to the community at large that here's someone not to be trifled with.

So something like law and order are maintained, and the status of the leadership preserved, through the threat (and, if necessary, the reality) of violence--but also through patronage delivered to constituents, and through satisfying the widespread craving to have

a hero to admire, who can tell you what to do, especially when there's a threat from outside the group.

Male chimps are obsessively motivated to work their way up the dominance ladder. [In the discussion omitted earlier, they say that as with humans this varies, some are ambitious, some content with their lot.] This involves courage, fighting ability, often size, and always real skill in ward-heeler politics. The higher his rank, the fewer the attacks on him by other males and the more gratifying instances of deference an submission. But the higher his rank, the more he will be obliged to take pains to reassure subordinates.

The alpha male, merely by virtue of his exalted status, inspires conspiracies to depose him. A lower-ranking male may challenge the alpha by bluff, intimidation or real combat, as a step toward reversing their relative status. Especially under crowded conditions, females play a central role in encouraging and helping to implement coups d'Etat. But the alpha male is often prepared single-handedly to take on coalitions of three, four or five opponents. Political assassination--that is, dominance combat in which the loser dies--is rare.

Any given fight is likely to stimulate other fights among unrelated or even unaffiliated parties. One combatant may poignantly appeal for aid from passers-by, who may, in any case, be attacked for no apparent reason. Everyone's hair stands on end. Perhaps longstanding resentments flare. General mayhem often results.

Alliances are made and broken. Loyalties shift. There is bravery and devotion, perfidy and betrayal. No dedication to liberty and equality is evident in chimpanzee politics, but machinery is purring to soften the more hardhearted tyrranies. The focus is on the balance of power.

In this complex, fluid social life, great benefits accrue to those skilled in discerning the interests, hopes, fears and feelings of others. The alliance strategy is opportunistic. Today's allies may be tomorrow's adversaries, and vice versa. The only constant is ambition and fixity of purpose.

Males have special reasons to avoid permanent rivalries. In hunting other animals and in patrols into enemy territory, they rely on one another. Mutual mistrust would be dangerous. Also, they need alliances to work their way up the promotion ladder or to maintain themselves in power. So, while males are much more aggressive than females, they also are much more highly motivated toward eventual reconciliation.

In zoo after zoo, males--especially high-ranking males--exhibit a degree of measured restraint under crowded conditions that would be unthinkable if they were free. Captive chimps are much more likely to share their food. Captivity somehow brings forth a more democratic spirit. When jammed together, chimps make an extra effort to get the social machinery to hum. In this remarkable transformation, it is the females who are the peacemakers. When, after a fight, two males are studiously ignoring one another as if they were too proud to apologize or make up, it is often a female who jollies them along and gets them interacting. She clears blocked channels of communication.

At a large chimp colony in a zoo in the Netherlands, every adult female was found to play a therapeutic role in communication and mediation among the petulant, rank-conscious, grudge-holding males. When real fights were about to break out and males began to arm themselves with rocks, the females gently removed the weapons, prying their fingers open. If the males rearmed themselves, the females disarmed them again. In the resolution of disputes and the avoidance of conflicts, females led the way

Chimpanzee females and their young have deep bonds of affection, while the adolescent and adult males seem more often mesmerized by rank and sex. The young revel in rough-and-tumble play together. Occasionally chimps of either sex will endanger themselves to help others, even those who are not close relatives. Male bonding on a hunt or patrol into enemy territory is palpable. Clearly there are opportunities for civil, affectionate, even altruistic behavior in chimpanzee society.

Females are not born knowing how to be competent mothers; they must be taught by example. The investment of time required of the mother is substantial: The young are not weaned until they're 5 or 6 years old, and they enter puberty around age 10. For much of the time until weaning, they're unable to care for themselves. They're very good, though, at clutching their mother's hair as they ride upside-down on her belly and chest. So long as they allow the infant to nurse whenever it wants, chimp mothers are usually infertile and unattractive to males. Without the males constantly hassling them for sex, they're able to spend much more time with the kids.

Chimp mothers use corporal punishment very rarely. Infant males learn the conventional modes of threat and coercion by closely observing older males, and they soon attempt to intimidate females. Before reaching adulthood, nearly every male has obtained submission from nearly every female. The youngsters yearn to be apprentices and acolytes of the older males, and are simultaneously nervous and submissive and hopeful in their presence. They're looking for heros to worship.

(From Carl Sagan and Ann Druyan, _Shadows of Forgotten Ancestors_, Random House, due out this fall. Excerpted in _Parade Magazine_ for June 7, 1992, issued with _The Boston Sunday Globe_.)

What perceptions might be being controlled here?

Bruce bn@bbn.com

Date: Fri Jun 19, 1992 4:55 pm PST
Subject: trendy scince, behavior & control

[From Rick Marken (920619 14:30)]

Lance Norskog (920618)

Welcome! You have touched on my favorite subject. I completely agree that PCT has a tough time getting much attention because, as you say,

>Control systems are not the latest spiffy gizmo, therefore
>the brain can't possibly work like one.

But I think that those in the AI biz (as well as others who are interested in modeling aspects of human and animal behavior) have another, even more fundemental, problem that leads them to ignore PCT. The problem is that they don't know what PCT is trying to explain. What PCT is trying to explain is CONTROL (or purposive behavior). Control is

NOT what AI, neural nets, expert systems, subsumption architectures, Beer Bugs, attractor models, etc etc are trying to explain (although these models may end up doing some controlling by accident [from the point of view of the modeller]). The "AI type" models are attempts to imitate BEHAVIOR -- that is, they attempt to mimic the various outputs that are generated by an organism over time -- the "intelligent" looking outputs like playing chess, proving theorms, getting around obstacles, navigating, conversing, etc etc. The word "behavior" is, unfortunately, used to refer to both this "output generation" process as well as to "control". So PCTers and AI types often think they are talking about the same thing, when, in fact, they are only using the same word, "behavior".

I see the trendy "AI" type models as the modern incarnations of the the devices built in the 17th century that could perform impressive (for the time) life - like sequences of actions. AI software is an advance over these devices only in that it can generate even more impressive outputs. But the arcitecture of these modern systems is basically the same as the architecture of the old devices -- an input - output architecture. This is the kind of architecture that makes sense when one is trying to generate outputs; it's the wrong architecture for producing control.

So to me, the question is "why aren't behavior modellers -- including AI types -- willing to learn about control (not necessarily control theory)"?

I think the answer is clear from control theory itself; doing so would require seeing that their current goals, and the means they have learned to achieve them, are based on a misconception (that behavior is generated output rather than control). I can see why people would not be seriously interested in finding this out -- and I sympathize with these people (though they can sometimes be awfully irritating with the arrogance of their self-deceptions).

Best regards Rick

Date: Sat Jun 20, 1992 11:23 am PST
Subject: standards, conflicts

[From Rick Marken (920620)]

One of the things that has particularly irritated me about the current political dialog about values (the one going on in the outside world -- not on CSGnet) is that the people who are pushing "family values" most ardently are also the people who have most ardently pushed one of the most fundemental (and, I think, destructive) values of our (US) society -- the value of CONFLICT (also called COMPETITION).

Every red-blooded American knows that competition is what makes for successful economies. The basic idea (as pink-blooded little me understands it) is that consumers are like judges at a beauty contest (a uniquely American event itself). Producers (or goods and services) compete to win the patronage of the customers. This competition leads to better and better products from producers (in the sense that they are the products that best meet the customer's needs or wants).

This scenario has one little problem that only Americans with pink tainted blood might ever even deign to point out; in competition like this there is generally a winner and a loser. What happen's to the loser? America doesn't like losers so we ignore them or blame the loss on personal failings (not being a REAL MAN (or WOMAN)). Pinko types like me, however, don't think that losers are just valueless trash; they are worthwhile control systems, with intrinsic reference signals of their own. I worry about the losers because societies with lots of them around tend to be very precarious -- and have to take strong measures to make sure that the losers don't try to just take stuff from the winners.

I don't like the "value" attached to competition in this society. I like the "value" of cooperation and community. I think society's emphasis on the importance of "being # 1" or "fighting to get to the top" is far worse than the lack of emphasis on "family values" and the other bullshit being discussed in the media. But I doubt that Quayle and Bush will come out im favor of the value of "cooperation" and "community". Do I have bad standards? Is it wrong to dislike competition and to like cooperation?

I will admit that competition (conflict) can accelerate the development of technologies that might help the parties to the conflict "win". Thus, two companies making widgits might progress faster toward the goal of making the "best" widget (the one that satisfies the market best) because they are in conflict (they have to keep improving the widget -- the output of each system-- or lose the conflict -- have their market share of widgets become much lower than their reference).

I think it is this "good" result of competition that has impressed economists. But is this the only way to organize an economy that produces that widgets that we all need to control what we want to control? Must there be winners and losers in order to have an economy that meets the requirements of its members (the winners, anyway). Can't we organize a society in which everybody is a winner (can control what they need and want to control) -- and can't we do it without coersion (the approach that communism used?). It seems to me that the economies of some of the scandanavian and western european societies approach a nice compromise between capitalistic individualism and socialistic communalism. Why don't we learn from those economies?

Best regards Rick

Date: Sat Jun 20, 1992 5:37 pm PST
Subject: standards & conflect

[from Avery Andrews 920621.1125]
(Rick Marken 920620)

It sounds like you ought to be living down here, not up there.

It was just claimed in a newspaper column I read here that the formerly somewhat pinko countries that have gone furthest down the laissez-faire road (the UK & New Zealand) are the ones that are currently in the biggest economic mess.

Avery.Andrews@anu.edu.au

Date: Sat Jun 20, 1992 7:46 pm PST
Subject: Re: standards, conflicts

Hurrah! I agree violently :-) with the judgement that competition is often (maybe very often) misplaced and misapplied. I tried to teach my kids that 20 years ago. I'm still convinced that was the right thing to do even though sometimes it seems to put them at a severe
disadvantage in dealing with the yuppies and other "me generation" types.

(What's the metaphor for (no, that's not a stammer) using a keyboard to "shoot from the lip"?

Ray Allis

Date: Sun Jun 21, 1992 7:26 am PST
Subject: Sameness; chimps

[From Bill Powers (920618)]

Martin Taylor (920619,19)]

OK, I think I see the "sameness" problem now. I had thought initially, Martin, that there were 4 objects of one kind in the midst of n objects of other kinds, and I didn't realize that _continuous_ tracking was required (i.e., you can't break for lunch and then go back to noticing the 4 objects). I didn't realize that you were selecting 4 objects for tracking out of many other objects OF THE SAME KIND. All this changes the nature of the problem rather drastically.

As you describe the problem, some number of objects, indistinguishable from
each other except by position, is presented to the observer. One to four of
these objects are pointed out to the observer as individuals to be followed. The objects then all begin to move about in some manner at some speed, and when they stop you ask the observer to point out the original one to four objects.

The question you ask, as I now understand it, is how the observer knows that these are the "same" objects. And my answer is that he doesn't. There is in fact no way to verify that they are -- that during the melee, one of the non-indicated objects didn't suddenly swap places with a tracked one. The initial condition is that there are no distinguishing characteristics other than position. If this remains true throughout, it is true at the end as well.

There is, of course, an hypothesis that they are the same individuals. But if, at any point during the moving-about, one of the non-indicated individuals momentarily coincided with an indicated one (dots or other identical figures on a computer screen), no observer would be able to differentiate them, even by position, at that moment. So the two individuals that depart from that position afterward have lost their distinct identities. The same result would occur if the two individuals moved in smooth paths, or in paths with a discontinuous kink in them at the
point of coincidence. Only the hypothesis that they move continuously would
lead to choosing the one set of identities after the coincidence, and rejecting the other. So the question is how that could be done.

The problem you state breaks down into two questions: first, how does a person keep track of SINGLE individual in a set of identical moving objects, and second, given that the first question is answered, how can more than one such process occur in the brain at the same time?

If the objects are all alike, then clearly their "identity" in terms of their classification at the beginning, before movement starts, is irrelevant. The objects would all be classified alike. They would, furthermore, all have the same configuration, the same sensations, the same
intensities. They would all be showing the same transitions (none). One might initially define them in terms of spatial relationships, but as soon as the objects began milling about, all initial relationships would be destroyed. If the motions were random, no hypothesis about a distinguishing
relationship (such as "the top center object" or "the object between two others just to the left and right of it") would be borne out. Only the hypothesis of continuity of path could possibly distinguish any object from
another if their paths led to a momentary coincidence. But if the objects moved in curves that met at a tangent, continuity of first derivatives would no longer suffice to distinguish them; if the period of coincidence included a brief straight-line segment, no number of derivatives would suffice. Identity would be lost, and any further identification of the original object would be only a guess.

Following a given object successfully, then, requires that no two objects ever coincide in a way that renders all hypotheses about their characteristics irrelevant. In practical terms it probably means that the objects must never coincide at all for any discernible length of time.

For a single object, we are left, then, with simple x-y spatial tracking. A
gate of finite size must be placed about the image of the object (neural image), small enough that normally only a single object occupies it. Movements of the object within this gate are detected, and the position of the gate is controlled to keep the object at the center. The control parameters of the gate are such as to be able to follow changes in the

direction and velocity of the object. Whatever object is in the gate at the
end of a period of continuous tracking is understood to be the "same" object that was there initially (although, under the conditions you state, there is no way to verify this). If tracking is interrupted for some time, all indication of identify is lost.

To distinguish objects that momentarily coincide, in the absence of any distinguishing features in the images of the objects, certain assumptions must be built into the tracking circuitry. For example, it may be that objects never move in a curve with less than a certain radius (no true right-angle movements, instantaneous reversals, etc.). The bandwidth of the
tracking system would then be limited so that an object that too-suddenly changes direction passes out of the control range. Thus if two objects crossed paths and both appeared in the tracking gate, the average position would suddenly depart from the smooth path being followed, but the tracking
gate would not be able to follow immediately this sudden change of direction. The gate would be momentarily disturbed, but not enough to "lose
track" of the object originally being followed; as the other object passed out of the tracking range, the tracking gate would quickly center on the original object again, its "momentum" having carried it more or less in the
original direction for a while during the disturbance. This "momentum", of course, is the result of the deliberate restriction on the speed with which
the gate can change its vector velocity. [Pat Williams, when you get back on the net: note the relevance of this to your problem of changing bit-maps
to line drawings].

If the second object remained within the gate for too long a time, the tracking system, which can work only in terms of centroids, would be unable
to pick the right object when the two (or more) objects diverge again. So the size of the gate matters as much as the dynamic characteristics of the gate movements. The more likely it is that two objects can occupy a given size of gate area, the smaller the gate must be to retain a good probability of following the "same" object. If the gate is too small, however, tracking even one object would become difficult.

The gate would be entirely neural in nature. It would correspond to focussing attention on a particular small but movable area in the visual map of the outside world.

Now: the second problem, which is "how many gates of this kind can be maintained at the same time?" Apparently, the answer is three, or sometimes
four. But I would think that this would depend on the characteristics of motion and the density of objects in x and y. If the motions entailed abrupt and arbitrary changes in direction, and the field of view were densely-enough occupied, tracking even one individual correctly would become unlikely. So there are many parameters in this situation. The

apparent limit of four objects could be simply a matter of probabilities: the more objects that are being tracked, the greater the probability that at least one tracking-gate system would encompass two objects and make a mistake. I would set the limit at four only if three objects could ALWAYS be tracked without error and five objects could NEVER be tracked without error.

This long treatise, of course, is simply an exercise in modeling. A gated tracking system is a model of one way (the only way I can think of that would work) to follow a single object moving among identical-appearing objects. This model introduces certain parameters of the situation: characteristics of the tracking circuitry, which we guess at, and certain characteristics of the experimental situation that become relevant under that model. The model can be tested by seeing whether it predicts tracking success and failure under conditions where the model would succeed or fail;
the model can be trimmed up by varying the experimental conditions to see whether predicted effects occur, such as effects of object coincidences and
near-coincidences. In fact, it seems to me that it would be possible to measure the size of the gate and the parameters of the gate position control system in this way -- provided that the results continued to be consistent with the model.

Under this model, "identity" is irrelevant, in the sense of one object having any unique distinguishing characteristic at any level of perception.
And even "continuity" is, under certain conditions of near-coincidence in space, problematical.

Bruce Nevin (920419.1348) --

Sorry to be obtuse and to seem rejecting (or perhaps I should interchange "to be" and "seem"). But you got my true reaction, which was that this mode
of description isn't very useful, even as a point at which to start guessing at controlled variables:

> The alpha male is sitting bolt upright, jaw set, staring confidently into middle distance. The hair on his head, shoulders and back is standing on end, which gives him an even more imposing aspect. Before him crouches a subordinate, in a bow so deep that his gaze must be fixed on the few tufts of grass directly before him.

This observation, whether made of chimpanzees or humans, puts more imagination than observation into the picture. "Upright" is one thing; "bolt upright" is another, implying an imagined way of getting into that position. "Jaw set" is completely imaginary unless you can feel the efforts
(if any) involved in holding the jaw in that position. "Staring" is OK, but
"confidently?" "Into the MIDDLE distance?" Those observations tell us a lot
about the observer but nothing (verifiable) about the observed. Hair does stand on end, in certain places, but in a human being it is not a

particularly striking effect. Whether it creates an imposing aspect depends
entirely on how much you feel imposed upon by it. A "crouching" figure I can imagine, but calling this a "bow" is certainly going too far. And "subordination" is surely an interpretation, as is one's guess as to what part of the ground is being gazed at, and what it is within the field of view to which the crouching figure's attention is turned. My own guess is that this subordinate is focussing on a imagined scene of catching this threatening son-of-a-bitch alone with his back turned.

The colorful descriptions by Sagan and Drayan are meant to convey a message: that chimpanzees behave just like human beings in certain basic social respects. To make the message stronger, they use a lot of imagery deliberately designed to evoke familiar perceptions in the (human) listener. But in doing so, they bring in a great deal that isn't actually observed. Even worse, they reify subjective impressions, even when referring to human beings, making it appear that these impressions correspond to something objective in the world outside the observer (and incidentally attempting to qualify themselves as unbiased objective observers by the use of these biased and subjective descriptions).

I think this kind of rhetoric is all but useless as a way of understanding behavior. It's all on the surface -- literally, it's superficial. All it does is describe, and the description is so strongly biased by underlying concepts of cause and effect that it unconsciously pushes those concepts on
us. This way of describing nature focusses entirely on actions, outputs, side-effects. It scarcely touches on what such actions accomplish -- or even whether the effects noticed have anything to do with what the actors are trying to accomplish.

If control theory teaches us anything, it's that actions and outputs and side-effects thereof are only an indirect indication of what is actually going on. We see what the actors are doing to their bodies and the world around them, but we don't see what perceptions are important to those actors, or what states of those perceptions the actors are trying to achieve. We see, under the old way of thinking, that all these actions are causing certain effects in the situation. What we don't realize is that the
effects are what are causing, calling for, the actions.

The "perks" of the alpha male are simply what the alpha male, and probably any other male, wants. What are those things? Should we take it for granted
that some abstract condition called "dominance" is an end in itself, or that it even has objective existence? Is dominance something sought, or is it simply a side effect of being the strongest control system in a conflict
situation? Are the marks of submission effects of the dominance, or are they simply all that is left to do by way of control when alternative modes
of action have been defeated by superior force? Do "perks" entail obligations, or is it that when one set of goals has been accomplished, others come to the fore?

(to everybody)

I've ranted for years to the CSG that if we want to have a revolution, we must revolt. We can't just go on using the same old customary modes of observation, description, and explanation if we want to find the significance of the first new concept of human nature since Descartes. If people are going to try to make a smooth transition into control theory, preserving everything they had thought important up to that point and simply adding a few new interpretations, where convenient and supportive of
former beliefs, we are going to get exactly nowhere. Control theory gives us the chance to tear all of our old ideas down to their components and put
them back together into a new structure of understanding. There is a great reluctance even in the smartest people I know, many of whom are on this net, to give up on the old approach and really try out the new one. Everyone has something (and for different people, different things) that is
too valuable or true to give up or rethink. Everyone has past accomplishments that they don't want to analyze too deeply in terms of control theory, lest a flaw be found. That's just controlling for being right, and is quite natural.

But anyone who wants to be a control theorist has to start trying out things that seem unnatural, doubting what seems right, giving up what seems
valuable. We have to have faith that by such acts of internal destruction, we will arrive at something closer to the truth, and salvage what is really
worth salvaging, when the reconstruction, under new management, begins. God, I really sound like a mindless revolutionary. But this is the way it has to be. Otherwise we're just fooling around and trying to impress each other for our own entertainment.

Best to all, Bill P.

Date: Sun Jun 21, 1992 8:54 am PST
Subject: Dilemmas of competition and cooperation

[From Kent McClelland (920621)]

Rick Marken (920620)

Although I agree with and indeed applaud your sentiments favoring cooperation
over competition, I wonder whether you're making the choice sound a little too simple. An interesting book by Michael Billig and associates (Ideological Dilemmas: A Social Psychology of Everyday Thinking. London: Sage, 1988) has convinced me that such things are not a matter of either/or,
at least not in our usual modes of thinking. Billig et al. trace the history
of Enlightenment thought and show how contradictory values are built into the

public discourse on such issues. Racists, for instance, will typically preface their biased remarks with a disclaimer to the effect that they themselves aren't prejudiced against blacks but you really can't get away from the fact that . . . etc., etc. I have no doubt that Bush and Quayle could come up with many heart-warming remarks about the value of community and not see any contradiction at all between that sort of rhetoric and their
views on competition.

From an HPCT point of view, I think the question is how stable system concepts can come to be constructed from an amalgamation of values or principles that are often contradictory in practice. But maybe such mental
and moral flexibility is necessary for us to maintain the perception that the
world we observe is consistent with our preferred system concepts. As I believe Rick pointed out in the discussion of values on the net a few weeks
ago, a control system that was stuck with a single reference signal for a principle like honesty (or cooperation!) would be unable to vary its outputs
to maintain control of perceptions of the next higher level, just like an arbitrary restriction to a single setting for arm position would cripple your
physical control of bodily movements.

Kent

PS to Rick: You asked sometime recently for confirmations on whether people
had received copies of your revised Elephant paper. I got it. Thanks. Sorry I didn't confirm at the time.

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

Date: Sun Jun 21, 1992 11:59 am PST
Subject: Cooperation, Conflict, & System concepts

[From Bill Powers (920621.1200)]

Kent McClelland (920621) and Rick Marken (920620 & prev) --

Cooperation and conflict are outcomes of a social interaction. If people's goals are aligned, there will be cooperation or at least non-interference. If they are not, there will be conflict and competition.

Competition arises in our society as a consequence of system concepts and principles. One of these concepts has to do with position in a social hierarchy. The idea of the superior person, with others being inferior, sets the stage in some people for a desire to be, or be acknowledged as being, at the top of this social hierarchy. As achievement of this goal

requires a relative ranking of people, it is impossible for everyone in the
society to achieve it. If even two people wish to be perceived as number 1,
a conflict must arise because by definition only one person can be number 1
(or number anything). The existence of number 1 creates number 2: number n implies number n - 1. If one person wants to become a leader, followers must be found, and others who also want to be leader must be fended off, undermined, or otherwise prevented from succeeding. The striving for social
position is a pernicious ill in our society, which accounts for a great many of its problems.

An anecdote. At the newspaper where I once worked (1979-1990), both the head programmer and I became potential candidates (in the eyes of management) for head of the department of technical services at a pleasantly elevated salary. My "rival" and I had exactly the same attitude:
who wants to spend the day going to meetings, making out budgets, dealing with public relations, and chastizing people for breaking the rules? So we both turned it down. Management insisted. We persuaded them to try out an old hand at the newspaper in our department, and after a few weeks he resigned and went back to being a competent technician. The pressure continued.

Finally, we got together, all the refuseniks, and consulted with the man who headed up the parts department, doing the ordering and filing everything (mostly erroneously). He had that lowly position because the technical schools he had attended didn't teach him much that would justify his title of technician, he had a lot of seniority, and he was willing to do things the way we told him they needed to be done. We asked him how he would like a promotion to head of the department. He thought that was a wonderful idea -- surprising, but wonderful.

So we put up his name and helped him write his resume and job description, and management informed us that he would now be the new director of technical services at about $10k more than his former salary. Our man V. moved into a big new office with three windows overlooking the Chicago River, with a rug on the floor, and with 23 peons working cheerfully under him, including the head programmer and me. We ripped out a wall to make the
office even bigger and put in a conference table with fancy visitor's chairs. We got him his own desktop computer and taught him how to write with a word-processor, and then for about a year taught him how to read and
write, period. We set up Lotus programs for him to do the budget with. We called him boss, and always deferred to him and tried to make him look good. Whenever any difficult decision arose, he would come to us, and we'd confer about it and tell him what to say at the weekly management meeting. If the department needed something, or the paper needed something in our line, we'd write it up and submit it, through V., to management (as V.'s idea). When management wanted something done, V. would relay it to us and if we didn't think of reasons why doing it would be disastrous, we'd get on

it and do it.

Now the interesting thing was that as the years went by, V. got better and better at his job. He wrote better. He read better. He showed courage and principle. He made sensible decisions. In short, he became a very competent
manager -- and never once thought that he had control of anyone else in his
department. Management was happy with V. V. was happy with himself and us. And we were happiest of all, because we could go on doing exactly what interested us the most, exactly what we knew was best for the newspaper, and didn't have to do a job we all would have hated.

I don't mean to imply that everything in this department was perfect. There
were all the usual problems and conflicts, including conflicts with management outside the department. Many interpersonal problems existed and continued -- the usual. But the one problem we didn't have within the department was that of competition for social superiority. The problem simply didn't exist. Nobody wanted to be number 1 except V, and he knew it didn't mean a damned thing to the rest of us.

In fact there was almost no competition even professionally. It was decided
when the department started, under a long-gone department head, that there would be no shame in asking help; only in failing to ask for it and bluffing through and screwing up the newspaper. Ignorance was expected, and
mutual teaching and learning was expected. Even mistakes were expected and everybody, even people awakened at three in the morning, converged on the problem and helped to fix it. They didn't have to. Nothing was said if they
didn't. But they did. Trouble attracted people like flies. The job was to keep the newspaper running 24 hours a day, so well that all those users out
there wouldn't even know we were there. Everyone signed onto this goal and made it personal.

The average down time of our 17 PDP-11's and the rest of the system -- terminals, disk drives, printers, modems, typesetters, and other machinery that did all the word processing and typesetting of the news and classified
ads in the newspaper and kept track of the press runs -- was about 10 minutes. And this with an ATEX system that corrupted the data on a disk at least twice a day due to several unfixable bugs in the proprietary operating system.

You'll notice that I haven't even mentioned upper management. My newspaper got in deep financial trouble through two successive leveraged buyouts. It got a labor lawyer in as publisher and became a union-buster (our department voted down the idea of organizing). If the newspaper survives, it will be in spite of upper management. Upper management ran the paper strictly on the basis of social position and competition, same as most other businesses. The real newspaper ran itself, while upper management

(and the owners) did their best to flush it down the toilet while they vied
with each other for personal advancement.

There was a time, at one CSG meeting, when an old-timer in the control theory business came to a meeting for the first time and vented his spleen at one William T. Powers for leading the group in the wrong direction. He wanted to lead it. I said, "O.K., Bob, you've got it: go ahead." That was the last or second-to-last meeting he attended. What he didn't realize was that I don't lead the CSG. I don't want any power over anyone. In fact people who understand control theory just don't lead worth a damn. There's a lot of cooperation going on, but the competition level hovers around zero. When I start telling people in this group what to do, they say "Yes, Bill," and usually go right on doing as they please. That pisses me off, but I'm proud of it, too. It shows that they understand the message. Sometimes they have to remind me of it.

I've heard all the arguments in favor of competition. I don't believe them.
I don't think that people with contradictory goals accomplish anything but building up their muscles and cancelling the effects of someone else's muscles, leaving little effort available for real progress. I don't believe
there is a "top" in the social hierarchy -- I don't even believe there IS a
social hierarchy. And as long as I don't believe that, there is no social hierarchy for me. This doesn't endear me to people who want such a hierarchy to exist, but that's their problem. There's nothing I want from anyone that would make it worth while to play that game. Not even the privilege of living.

And I know for certain that when, in some microsociety, people manage to do
without this concept of Number One, everything magically works better: shared goals are accomplished smoothly, easily, and with great pleasure. People get smarter, because they aren't wasting their time and effort trying to counteract what someone else is doing.

I haven't got this system concept worked out in any detail -- talking about
it too much tends to reduce it to procedures and slogans, anyway. But what I do understand of it, I want to sell. It defines the kind of world that I find worth living in. All I can do to create that world is to persuade others who will persuade others that it's worth a try.

"Any new idea, Mahound,is asked two questions. The first is asked when it's
weak: WHAT KIND OF AN IDEA ARE YOU? Are you the kind that compromises, does
deals, accomodates itself to society, aims to find a niche, to survive; or are you the cussed, bloody-minded, ramrod-backed type of damnfool notion that would rather break than sway with the breeze? -- The kind that will almost certainly, ninety-nine times out of a hundred, be smashed to bits; but, the hundredth time, will change the world.

'What's the second question?' Gibreel asked aloud.

'Answer the first one first'."

Salman Rushdie, _Satanic Verses_, p. 335.

And thanks to Mary for finding it.

Best, Bill P.

Date: Sun Jun 21, 1992 12:31 pm PST
From: Dag Forssell / MCI ID: 474-2580
Subject: Belief Systems, misc

[From Dag Forssell (920621-1)

I should have posted this "reinforcement" of Bill's point in his post on Belief systems long ago. Better late than never?

Bill Powers (920429.0900)

>At the level of systematic belief, both principles and reasoning become
>subservient to preservation of the belief system. When you look at the
>arguments against purposiveness in behavior that were advanced -- and
>thought rather clever -- in the early parts of this century, you find
>elementary logical errors and straw-man arguments that wouldn't convince
>a schoolchild if the subject were something else. You find abandonment
>of principles of scientific detachment and objective argument in favor
>of emotional attacks and innuendo. The belief system justifies these
>alternative uses of principle and reason, because above all, the belief
>has to remain true. WHEN YOU ARE DEFENDING SOMETHING THAT IS ABOVE LOGIC
>AND PRINCIPLE, LOGIC AND PRINCIPLE MUST BE BENT TO THE HIGHER PURPOSE.

(CAPS emphasis by Dag)

Editorial pages Los Angeles Times, May 8, 1992:

THE JURY'S THINKING HAS BEEN HEARD BEFORE

Verdict: Police footprints on the victim's face couldn't persuade a Miami panel.

By ANDY COURT

As I listened to a juror explain that Rodney King was in "control" during his beating by Los Angeles police officers, I thought of Bernie and Rubina and Bill, down in Miami. They were nice people, and they, too, reached a verdict that set parts of a city on fire.

What they told me more than a year ago is relevant now because it might dispel the illusion that most of us still embrace: that the King verdict was the work of fools or overt racists. Something much more universal is at work, and race, in my opinion, is only one part of it.

Bernie, Bill, Rubina and nine others served as jurors in a federal civil-rights case against six Miami narcotics officers. The allegedly brutal officers represented a rainbow coalition of blacks, whites and Latinos; the victim was Latino. The jury, though mostly white, included three blacks and one Latino.

The prosecutors didn't have a videotape this time, but they had just about everything else. Leonardo Mercado, a smalltime drug dealer, had been beaten to death after entering a house with the officers. His corpse had 44 bruised areas, and marks on his forehead corresponded to some of the officers' sneaker-prints. A patrolwoman who did not participate in the beating testified that three of the defendants encouraged her to kick Mercado while he lay on the floor bleeding,

Nonetheless, the jury acquitted the officers of some charges and couldn't agree on the rest. After interviewing 11 of the 12 jurors, here's what I found:

Richard, a 38-year-old engine mechanic, said (during deliberations) that Mercado was "only a drug dealer, anyway."

Rubina, a 53-year-old saleswoman, didn't believe several prosecution witnesses from the neighborhood because "these are the people we're paying the policemen to protect us from."

Herbert, a 59-year-old airline mechanic, believed that "criminals give their civil rights away when they elect to lead a life of crime."

Bernie, a 48-year-old butcher, thought the police were guilty, but he changed his vote because "I didn't want to be the one that was sitting out there with them pointing at [me]."

Most telling, perhaps, was one juror's observation that the officers had to be found guilty "beyond an absolute doubt." This juror had single-handedly changed the standard of doubt in a criminal case. I suspect he did so because he felt more sympathy for police fighting the drug war than for a drug dealer with a violent past.

Most of these people weren't racists or fascists. In fact, they appeared so well-intentioned, so intent on applying the law as the judge had explained it to them, that it was all the more painful to witness how far they strayed from the realm of common sense.

They were working-class people who believed what the defense said about the defendants being the only thing standing between them and the chaos of the streets.

As one lawyer put it, most of the jurors had "never been on the wrong side of a nightstick." They did not sell drugs on street corners or engage in high-speed chases with police. Nor were they psychologically prepared to uphold the rights of those who did.

"To know what actually happened," one of the Miami jurors told me, "you'd have to be there or have a tape of it." Now it appears that even a tape isn't enough. That's because the problem is attitudinal. The jurors who produced the Rodney King verdict are a reflection of the American middle class's law-and-order mentality, which has been fired by the Administration's ill-conceived war on drugs and the widespread perception that too many' criminals get off on technicalities.

Convenient as it is, the bashing of the King jury is hypocritical, because a lot of Americans would have done the same misguided thing when the fate of these veteran police officers was put in their hands. In such situations, a weighing of souls occurs, and unless there are allegations of corruption, the police will almost always win over the criminal suspect.

The sad truth is that people not so different from ourselves as we'd like to believe will undertake Herculean feats of logic to acquit officers of blatantly brutal acts. They seem to sense that the police are, "us" and the criminal suspect is "them"-- and apparently "we" don't ever expect to end up on the wrong side of their nightsticks.

Footnote: Andy Court is editorial director of American Lawyer magazine, where material for this article first appeared.

Another newsclipping. Extract from Sunday book review in the Los Angeles Times a few weeks back. Sorry, no date left on torn page.

......."Who Will Tell The People" is one of the first books I have seen that talks at length about one of Washington's dirty little secrets: How much Senators and Representatives hate the folks they left behind. "Politicians are held in contempt by the public." writes Greider. "That is well known and not exactly new in American history. What is less well understood (and rarely talked about for the obvious reasons) is the deep contempt politicians have for the general public."

Exactly. Politicians, like the cops and emergency-room nurses Greider and I used to work with, tend to see people at their worst. The difference is that cops and nurses are there to help those people but politicians come to use them to help themselves. Laughing anecdotes about the naivete and grabbiness af constituents, those simpletons, is common conversation in Washington. And, human nature being what it is, each time polls show resentment of political pay and perks and plummeting public regard for the distinguished labors and sacrifices of members of Congress. those members remind each other that their constituents are still getting dumber every day...........

The reviewer notes that the author, Greider, has no solution to the problems he diagnoses.

I relate this to questions of business leadership. How many of the captains of industry feel contempt for the people who carry out their orders? How many leaders are of a mindset to recognize the individual autonomy of every person in their organization?

Competition VS. Cooperation:

Kosaku Yoshida: New Economic Principles in AmericaCompetition and Cooperation A Comparative Study of the U. S. and Japan

Abstract: The current decline in the U.S. economy has come about from the excessive practice of free competition. In previous centuries, unfettered competition and rugged individualism resulted in American economic prosperity. But the American environment has changed. Now, these factors which once created U.S. economic prosperity threaten to create the opposite. The author examines how the Japanese concept of cooperation can work with the Western principle of free competition to revitalize American competitiveness in the global market.

(The Columbia Journal of World Business Winter 1992 Volume XXVI, Number IV)

Post formats:

Sometimes, the length of individual lines do not fit on my screen. When that
happens, you get a lot of little orphans. I choose to reformat text to eliminate
these annoying, misplaced hard returns, and that takes a little time.

Of late, Bill's posts have a lot of these, but some others' too. To avoid this in my own posts, I set my word processor to courier (a fixed spacing font) in 11 points. (With another printer designated, I might set 10 characters per inch). (The margin setting matters, too). I save to disk in ASCII format, which places hard returns at all line endings, and post to MCImail and Internet from disk. I hope my posts have few enough characters per line, that it is easy to add > or >> in front without exceeding the screen line length.

Is this a generic problem or just mine? (I set 10 point text when I read and print, with one inch margins - perhaps I should change those settings. Now I have already written this. I'll send it).

Dag

Date: Sun Jun 21, 1992 12:36 pm PST
From: Dag Forssell / MCI ID: 474-2580
Subject: Promoting PCT

[From Dag Forssell (920621-2)

On June 5, I shared a letter (version six) to approach CEO's in American industry. 300 letters drew four replies in just over one week. (One from one of the few Japanese companies I addressed). I am delighted, since the response rate was greater than zero.

I am trying hard to capture the CEO's interest while giving an accurate impression of what I offer. Length of the letter is not a primary issue.

Holding attention and arousing curiosity is. I also do not want to put down those who make the effort to study the important subject of psychology, wether CEO's or psychologists, or both. (Quite the contrary)! It is the "science" that is inadequate.

Here is version eight. Since "everybody knows" what control is, I have adopted the term: "cybernetic control" to indicate something different and stimulate curiosity. Bold / underline has been replaced with CAPS. Comments and suggestions welcome.

(Purposeful Leadership TM letterhead)

Bill Powers, CEO June 21, 1992
CSG Forever, Inc.
73 Ridge Road CR 510
Durango, Co 81301

Dear Mr. Powers:

I am writing to introduce you personally to the first fundamentally new perspective on people that has been offered since 1637. Adopting it can mean improvements for your bottom line, productivity, quality and morale.

Costly people problems exist at all levels in American industry. Dr. W. Edwards Deming, pioneer in Quality Management, writes in "Out of the Crisis," page 85:

> "In my experience, people can face almost any problem except the problems of people. They can work long hours, face declining business, face loss of jobs, but not the problems of people. Faced with problems of people (management included), management, in my experience, go into a state of paralysis, taking refuge in formation of QC-Circles and groups for EI, EP, and QWL (Employee Involvement, Employee Participation, and Quality of Work Life).... There are of course pleasing exceptions, where the management understands... participates..."

There have always been natural leaders, successful salesmen, wise parents and good communicators. But it is rare that they can explain what they do and why. Their insight and skill seems intuitive. Some people in industry make the effort to master this subject. For most of us it takes extensive experience and attention to develop a consistently successful personal approach to dealing with people.

A fundamentally new perspective has been developed and is available for study. With it, understanding people does NOT have to be complex and confusing! The new perspective can be taught as an overview in a day and in considerable detail in three.

This new perspective gives an executive insight that allows him or her to inform, influence, align and lead people with mutual respect. S/he can teach people to be more effective and cooperative. Employees can be more satisfied, while the company as a whole responds better to the leader's direction and becomes more productive. The executive gains understanding and learns to function as well as those intuitively wise people. With practice even better, since s/he will have greater insight!

This perspective will also make it much easier to understand and teach Total Quality Management programs, such as the Deming Management Philosophy.

Describing this perspective so you get the point immediately is a Catch-22 challenge, because it is a different concept altogether from what predominates in our world today. Until you understand the principles, you cannot understand at all, and I need a few hours in class to convey the principles, before I can teach how to use them.

Over, please...

Let me use an illustrative analogy instead:

In an era when "everyone knew" that the earth was flat, scientific explanations were developed for navigation and astronomy. Many problems with those explanations persisted, but people worked around them. The explanations were taught to succeeding generations by experts. Non-experts took it all for granted without much thought.

I cannot say what "everyone knows" about human behavior, but experts on the subject employ a 17th century perspective of cause and effect to guide their research. Any book on experimental psychology tells you that the way to learn about behavior is to set up an experiment, then vary the stimulus (independent variable) and watch the response (dependent variable). With this scientific method our experts have done many experiments and formulated many explanations which have found their way into our language, culture and management practices. Non-experts take these explanations for granted without much thought.

Many problems with these explanations persist despite all the research, but people work around them. Our lack of consistent success indicates that we lack a good model or "paradigm" to help us understand why people do what they do. In our ignorance, we often spend our energies in debilitating conflict instead of in productive cooperation.

WHEN COPERNICUS AND THEN GALILEO INTRODUCED THE FUNDAMENTALLY NEW INSIGHT THAT THE EARTH IS ROUND (it has always been round), THE PROBLEMS OF NAVIGATION AND ASTRONOMY WERE PLACED IN A BRIGHT NEW LIGHT. The new insight did not invalidate the common sense observation that the earth appears flat locally, but science moved from a dead end to progress, which in a few centuries has brought us far.

But most experts of the old science could not comprehend the new paradigm, because they had already internalized the flat paradigm in all its details as their personal reality. With time, the experts died off, and new ones grew up, embracing the new paradigm on its merits because it solved many of those persistent problems. They internalized the new perspective, and science progressed from there.

Isaac Newton's "Principia Mathematica," published fifty years after Galileo, was resisted also for similar reasons. It took fifty years for it to be fully accepted. Looking back, we take it for granted. The evolution of science is much more than a steady accumulation of knowledge! 1 The process is creative. The opportunity for a revolution arises when a current paradigm fails to explain and competing paradigms are offered to provide better explanations. A struggle of many decades typically takes place, with the existing establishment continuing the development of the existing paradigm while outsiders and early converts champion a new one.

THE 20TH CENTURY UNDERSTANDING OF CYBERNETIC CONTROL APPLIED TO PEOPLE (people always control) PROVIDES A FUNDAMENTAL NEW INSIGHT THAT PUTS THE PROBLEMS THAT RESULT FROM HUMAN INTERACTIONS IN A BRIGHT NEW LIGHT.

Cybernetic control is as incomprehensible at first glance to a person trained in cause- effect thinking (which we all are to various degrees in our culture) as the idea that the earth is round was to a person trained in the details of a flat earth. The demonstration /test we offer shows this clearly. Still, an understanding of cybernetic control contains an explanation of the illusion of cause and effect in people, just like the understanding that the earth is round contains an explanation of the illusion of a flat earth.

Continued....

1 The phenomenon and process is described in Thomas Kuhn's seminal book: "The Structure of Scientific Revolutions," which introduced the term "paradigm."

Bill Powers June 21, 1992 Page 3

Another illustrative analogy is to say that we live in a maze where only the walls and passages are visible to us. The perspective of cybernetic control allows us to rise above the maze and see the structure. We can then set and reach our goal much easier.

The new perspective does not invalidate any wise common sense observation or practice. It just provides an enhanced understanding of seemingly intractable problems. It provides new diagnostic tools and shows why cookbook rules for behavior (programs which tell you what to do under certain circumstances) are inappropriate.

This perspective on cybernetic control in people is already well developed. But no doubt it will take time - well into the 21st century - before this successful breakthrough is embraced by a majority of experts. You can take advantage of what "everyone will know" in the 21st century right now to improve your company's competitive position. But because it breaks new ground, you must be willing to think for yourself to do it. You will actively participate in a scientific revolution when you adopt it.

The Purposeful LeadershipTM programs explain and translate this new perspective into skillful use of diagnostic tools that give you the capability to work on productivity. That includes effective communication, teaching effectiveness, resolving conflict, supporting self-motivation in employees, team building, Total Quality Management, leadership insights, effective performance appraisals, effective selling concepts, and development of corporate and individual mission statements. The executive learns how to build confidence, build trust, develop caring relationships and reduce stress.

The basic principles can be taught in a day to any attentive person, who can also verify them. People trained in the "hard" sciences will appreciate the scientific approach and elegant simplicity of the program, and everyone will be able to begin applying the principles as soon as they understand the underlying model and have had some instruction and practice with applications.

Some people will think that we promise a new way to control other people. It is precisely the other way around. We show how people control themselves at all times. When you understand cybernetic control you can work with people, rather than get into conflict despite the best of intentions.

Besides a consuming interest in this new development, I have 25 years management experience in engineering, manufacturing, finance and marketing. My formal education includes an MBA from the University of Southern California and a Masters degree in Mechanical Engineering from Sweden.

I will be pleased to send you a brochure and a free introductory 39 minute audio tape with script and illustrations. It demonstrates the perspective and explains the benefits, applications, background and content of our programs. The demonstration/test allows you to determine if your associates can recognize control in action. (I bet they can't)!

When you receive the introduction, I think you will find the demonstration both enlightening and entertaining. Please feel free to share it with your technical, operations and sales managers at any level for their evaluation. This is a win/win program to greatly increase the understanding and effectiveness of anyone who deals with people.

Sincerely,

(Page 4)

INFORMATION REQUEST

~ Please send me the introduction with a brochure.
  - ~ 39 minute audio cassette with complete script, including demonstration.
  - ~ Reading and illustration material only, no audio cassette.

~ Comment:

____
____
____
____

Bill Powers, CEO
CSG Forever, Inc.
73 Ridge Road CR 510
Durango, Co 81301

Correction:

Name: ____
Address: ____
____
____
Phone: (_____)__________________Fax:(____)____________________

Mail or fax this page with your request, or call direct.

Purposeful LeadershipTM
Dag C. Forssell
23903 Via Flamenco,
Valencia, CA 91355-2808
Phone: (805) 254-1195
Fax: (805) 254-7956

Date: Sun Jun 21, 1992 2:20 pm PST
From: Dag Forssell / MCI ID: 474-2580
Subject: Revolution, Leadership

[From Dag Forssell (920621-3)

Post formats:

I just downloaded two marvelous posts from Bill. They both had the formatting
problem. It went away when I adjusted my margin out another .3 inches. Sorry,
it was my problem, not a generic one.

Bill's date below is wrong?! Bill Powers (920618)

>I've ranted for years to the CSG that if we want to have a revolution, we
>must revolt. We can't just go on using the same old customary modes of
>observation, description, and explanation if we want to find the

>significance of the first new concept of human nature since Descartes. If
>people are going to try to make a smooth transition into control theory,
>preserving everything they had thought important up to that point and
>simply adding a few new interpretations, where convenient and supportive of
>former beliefs, we are going to get exactly nowhere. Control theory gives
>us the chance to tear all of our old ideas down to their components and put
>them back together into a new structure of understanding.

And in "An agenda for the control systems group 1986" (LCSII, page 171) you point out that the "human pie" has already been sliced.

I am working on an illustration of this. I think people need to be told specifically what it is about our daily language that is 1700's thinking.

As a tentative list, I have:

Mental illness
Depression
Addiction
Phobias
Anxieties
Compulsion
Reinforcement
Conditioning
Marital problems
Crime
Character
Preference
Aptitude
Personality
Intelligence
Self-esteem

Some of these are mentioned by Bill. Some I am not sure belong on the list.
Suggestions and comments are solicited.

Rick will note that I put Character on the list. I have accepted that: "It is ALL control." Perception and control is all there is!

Bill Powers (920621.1200)

>What he didn't realize was that I don't lead the CSG. I don't want any power
>over anyone. In fact people who understand control theory just don't lead
>worth a damn. There's a lot of cooperation going on, but the competition
>level hovers around zero.

In my book you lead this group, and well. You use what I have labeled Purposeful LeadershipTM.

In his book: "Leadership is an art," Max De Pree writes: "The first

responsibility of a leader is to define reality. The second is to say thank
you. In between, the leader is a servant."

You lead by offering information at the highest level. A new systems concept
of reality called PCT. You answer questions, explain and serve without letup.
You allow those who choose to follow to derive principles from the systems
concept information you offer (and so on down the entire hierarchy.)

You do NOT lead by coercion, threats and "rewards," you abhor those, but that
does NOT mean that you do not lead.

I for one am glad that you do lead!

Affectionately,

Dag

Date: Sun Jun 21, 1992 5:40 pm PST
Subject: "feedback guidance", conflict and cooperation

[From Rick Marken (920621)]

My "Hierarchical control of behavior" paper was sort of
rejected by a journal called "Consciousness and cognition".
The editor encouraged me to resubmit with more detailed explanation
of the model. I will try it -- the editor was very nice
about it. The interesting thing is that he had problems
finding anyone to review it; he managed to get one brief review,
which was basically positive -- with the usual misunderstandings
of PCT. The reviewer again brought up the evidence that
animals can behave even when deprived of feedback -- evidence being
the criminally evil "deafferentiation" studies (these would not
seem so horrendously evil to me if they weren't being done by
people who have NO understanding of how to study control systems).
The reviewer pointed me to a book that is, indeed, relevant to
the topic of my paper -- it's called "The organization of
perception and action" (1987) by D. G. MacKay (a UCLA linguist, not
the Christian neurophysiologist). What is interesting about this
book is that it tries to model the relationship between perception
and action in terms of the ol' cause effect model. Of course,
MacKay's model (a node excitation thing) is pretty silly but there
are some fun parts of the book. For example, he has a whole
section which explains why people are not interested in feedback
control models of behavior anymore; the reason? Because
so much behavior can be done without "feedback guidance"
(of course, he assumes that control systems work in a cause-
effect manner -- feedback causes behavior). One example he gave
is of studies that show that people can still speak even
after they have gone suddenly deaf. He didn't go into much detail
on this. Does anybody know about this evidence? Does the suddenly deaf

person speak the same as before going deaf? Do they immediately speak clearly (if they ever do)? Of course, MacKay has no idea that the occurrance of similar appearing "outputs" after sensory loss is not evidence that "behavior" can occur without feedback. What he has to show (to get me to abandon PCT) is that people can CONTROL after loss of the ability to sense a controlled variable. Thus, he has to start by showing what variables are controlled during speech (by doing the test); then he has to show that loss of the ability to sense all or some of those variables does not cause loss (or reduction) of control.

Nearly all studies which show "behavior without feedback" are based on the idea that "behavior" is output. If they understood that behavior is control, they wouldn't come to such silly conclusions. But the damage is done -- most psychologists assume that people can "behave" without feedback (just like those mechanical dolls of the 17th century). It is the accepted dogma, so it will be tough to get psychologists to reconsider a question that they think is already answered. Ah well. Perhaps some research directly related to these "behavior without feedback" studies has to be done and published. Anyone volunteer?

One interesting study that MacKay mentioned, which stongly points to the importance of feedback (though he didn't see it that way), showed that adapting a person with repeated presentation of the sounds /pi/ or /ti/ lead to a decrease in voice onset time (VOT) when these perceptually adapted people were asked to say /pi/ or /ti/. The study was done by Cooper and Naper and reported in JASA (Journal of the Acoustical Society of America) in 1975. Any linguists know about this? It sounds like there is a change in output (VOT) in order to produce an intended perceptual result (/pi/ or /ti/) via an adapted sensory input system. Is this interpretation reasonablee? It would have been nice if the study had been done quantitatively -- maybe it was. For example, the degree of change in VOT might be expected to depend on the degree of adaptation.

Dag Forssell (920621-3)

Great list.

>In my book you [WTP] lead this group, and well. You use what I have labeled
>Purposeful LeadershipTM.

If what Bill does is Purposeful LeadershipTM then, boy, am I for it!!
I just hope it can be taught.

Best regards Rick

Date: Mon Jun 22, 1992 6:37 am PST
Subject: looking at looking at chimps

[From: Bruce Nevin (Mon 920422 08:52:59)]

I do not claim that Sagan and Drayan have correctly described the perceptions that chimpanzees are controlling. I agree that there is no basis, either in their writing or in the reports of ethological observations which they summarize, for determining whether they have identified and described the relevant perceptions or not. (The "or not" is of course important.)

I do claim, and I think that you would agree, that projection of this kind is typical of what people do. One looks at what another person (or chimpanzee, or other creature) is doing. One imagines what one would have to do to accomplish those behavioral outputs. One imagines what it would feel like to do those things. One imagines what other aspects of one's perceptual universe would be like were one having those feelings, and doing those that one would do to produce like behavioral outputs-- one imagines what perceptions one would be controlling. One does this imagining with respect to the observed behavior of each of the players in a social interaction such as those Sagan and Drayan describe. For each, one imagines what it would take and how it would feel for one to do the observed things in the context of what one has imagined for the other players' actions. I think it is not controversial to say that this sort of "projection" is everyday fare for humans.

However, instead of attempting to eliminate this process as an unwanted interference with "objectivity," I would embrace it as being itself a crucial datum about our perceptions and our control of them, and an invaluable tool for insight into them; this precisely because we all do it, and because we know (or are confident) that we all do it, and because we have acted on this assumption all or most of our lives, evidently in concert with others so doing.

In particular, I make the further claim that the association of particular manners of behavior (comprising behavioral outputs), emotional states, and social roles (participation in mutually recognized social relations) is learned and indeed taught as part of how to be an adult member of one's society.

>This observation, whether made of chimpanzees or humans, puts more
>imagination than observation into the picture. "Upright" is one thing;
>"bolt upright" is another, implying an imagined way of getting into that
>position. "Jaw set" is completely imaginary unless you can feel the efforts
>(if any) involved in holding the jaw in that position. "Staring" is OK, but
>"confidently?" "Into the MIDDLE distance?" Those observations tell us a lot
>about the observer but nothing (verifiable) about the observed.

Exactly so.

What would tell us something about the observed alpha male is that these various enactments of assymetrical social relation (as we imagine them to be, as they would be if we were engaging in them with other humans) are indeed asymmetrical. That only one male is in the "superior" role in transactions with all others in the group, though others are recognizably so in relation to some others; that if the alpha male is deposed (as we imagine it), he is no does all these things in the "superior" role wrt all others, but rather the one deposing does; that young males learn the ways of enacting superior and inferior social roles and practice them. What all this (and more) tells us by its consistency and pervasiveness is that the alpha male and the other members of the group are controlling some perceptions which may be like those that humans control when they do analogous things. What those perceptions are, their reference levels, their relations to other perceptions, etc., all this is surely to be determined.

>Those observations tell us a lot
>about the observer but nothing (verifiable) about the observed.

So what do they tell us about the observer? That is of great interest to us. The same processes of observing behavioral outputs and imagining how one would experience a perceptual universe in which one produced like outputs oneself is at the heart of much of communication, which is about social relationships. One imagines (and remembers) the experience of enacting both (or each) of the roles that one perceives being enacted.

Bateson lays out the issues of (nonverbal) communication in many discussions, for example, in "Problems in Cetacean and Other Mammalian Communication" (reprinted in _Steps to and Ecology of Mind_ p. 364 ff.), where the discussion of the wolf pack is particularly apt. I will not type it, as it is accessible and as the effort will be ill timed if the above suggestions are controversial after all.

Bruce bn@bbn.com

Date: Mon Jun 22, 1992 6:58 am PST
From: marken
EMS: INTERNET / MCI ID: 376-5414
MBX: marken@aero.org

TO: * Dag Forssell / MCI ID: 474-2580
Subject: blindmen

Dag

Here is copy of latest version of my paper -- it's still rather terse. Comments appreciated.

Rick

The Blind Men and the Elephant:
Three Perspectives on the Phenomenon of Control

Richard S. Marken
June 9, 1992

Abstract - Psychologists have described behavior as 1) a response to stimulation 2) an output controlled by reinforcement contingencies and 3) an observable result of cognitive processes. It

seems like they are describing three different phenomena but they could be describing one phenomenon -- control -- from three different perspectives. Control is like the proverbial elephant studied by the three blind men; what one concludes about 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.

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 is a response to stimulation, an output controlled by reinforcement contingencies or an observable result of cognitive processes. Each of these ways of describing behavior is what would be expected if people were describing 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. These descriptions gave a true 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 a true picture of some aspects of behavior but a false picture of behavior as a whole. To see why this might be the case it is necessary to take a close look at what it means to control.

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 stable. Organisms 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

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

```
                     _______
d --- k.e ---> s --- | k.o | ---> r
               ^     _______      |
               |                  |
               |_____  k.f ______|
```

____

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:

(1) $$r = k.o\,(s^*-s)$$

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:

(2) $s = k.f\ (r) + k.e\ (d)$

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:

(3) $r = 1/((1/k.o)+k.f)\ s^* - k.e/((1/k.o)+ k.f)\ d$

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:

(4) $r = s^*/k.f - (k.e/k.f)\ d$

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. The organism does this by varying responses to compensate for variations in the environment 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. Equation (4) shows that behavior will look this way when the reference value for stimulus input 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 so-called "reflex" behavior, such as the pupillary response, where changes in a stimulus variable (such as illumination level) lead to changes in a response variable (such as pupil size). Of course, this relationship between stimulus and response is precisely that which is required to keep a sensory variable (sensed illumination) at a fixed reference value, s*.

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. Equation (4) shows that it would look like the feedback function controls 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; it is kept constant by maintaining the test animal at a fixed proportion of its normal body weight. The environmental variable, d, is the reinforcement, which, if it is food, is 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 term "ratio" refers to the number of responses required per reinforcement . So a "ratio 10" schedule is one in which the organism must make ten responses in order to get one reinforcement (a 10 to 1 ratio). This 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 complex 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. Cognitive behaviors are 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

(7) $$r = s^*/F + K$$

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 look at 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 behavior looks like different phenomena from different perspectives. 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, 1973). 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 seems like 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 input at reference values. Indeed, once you know the sensory inputs that are being controlled by the organism, all aspects of its behavior can be predicted from a knowledge of the laws that relate the organism to the environment. A controlled sensory input is called a controlled variable and s* is the reference value for a controlled variable. There are methods, based on control theory, that can be used to determine what sensory variables are being controlled by an organism 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

Marken, R. S. (1988) The nature of behavior: Control as fact and theory. Behavioral Science, 33, 196- 206.

Marken, R. S. (1992) Mind Readings: Experimental studies of purpose. GravelSwitch, KY: CSG Publishing

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

Staddon, J. E. R. (1979) Operant behavior as adaptation to constraint. Journal of Experimental Psychology: General, 108, 48- 67

Date: Mon Jun 22, 1992 8:09 am PST
Subject: Leaders and followers

[From Bill Powers (920622.0800)]

Dag Forssell (920621) and Rick Marken (920621) --

Thanks for the vote of confidence, but I'll stick to my definitions. I'm first of all an explorer of a new idea, second a teacher, and third a student. These are things that anyone can be. On this net, everyone plays these roles at one time or another. For one person to be any of these things does not rule out another person being the same things, at the same time, in the same group.

But "leadership," it seems to me, is a role in a social hierarchy. It requires followers. It opens the door to competition and conflict ("I can lead better than he can, so follow me and not him"). The worst result, from my point of view and in my circumstances, is that followers learn from a leader how to follow, not how to explore, teach, and learn.

The attitude of followers toward leaders, in my experience, often tends to be one of admiration, deference, blind loyalty, and even hero worship. It's the attitude of a child toward a favored adult. Many leaders like being on the receiving end of this attitude. It confers power, it allows the leader to indulge in egocentric thinking, it protects the leader from criticism and accountability. The leader can arrive too easily, with the connivance of the followers, at the idea that he or she makes fewer mistakes than ordinary people do. The leader can point to the support of the followers as a way of showing others, outside the group, that there must be something superior about the leader (so they would be better off becoming followers, too). Leaders are corrupted by their followers, and willingly.

I do have things to teach. I like teaching. I have a lot to learn. I like learning, too. But the main thing about me is that I was lucky enough to have the right, rather odd, combinations of knowledge and ignorance to make something new of an idea invented by others. This new idea is far more important than I am; it will be remembered long after I am forgotten. I want to work beside others who also understand this idea and think it's important, so the idea will go on living after I have lost the knack.

I don't want to think that when I disappear from the scene, my followers will cast around for someone else to follow, someone with another admirable idea, and will never think of carrying on for themselves what we are doing together now. But if they are followers, and don't think of themselves as independent explorers, as teachers with something to teach -- if they have learned only to follow -- how can they carry anything onward when they are cast adrift?

I'm too conscious of my own failings and ignorance to have any confidence of maintaining my integrity without others who will tell me when I'm making a fool of myself, or have misunderstood something, or have overlooked something, or am simply on the wrong track. Followers won't believe that I'm ignorant or have any failings; to tell me I'm a fool would make them fools, too; to tell me I've misunderstood would be to dare to have an independent opinion contrary to the guru's. To think that I've misunderstood something would be to make room for the thought that I have taught something wrong, maybe the whole thing -- after all, if you just take the word of the leader for everything, the only reason you have for believing anything is that the leader said it was true. If the leader is clearly wrong about something, the whole flimsy structure collapses.

I don't need followers. I need friends, people I love and who feel love for me, colleagues, equals. I don't want to be all alone on a pedestal.
So I don't want to be a leader, or a guru, or a saviour. I refuse. Just as I refuse, for all the same reasons, to be a follower.

Mary says that the second answer to the question, "What kind of idea are you?" is -- "How will you behave when you win? When your enemies are at your mercy and your power has become absolute, what then?"

Salman Rushdie, _Satanic Verses_, p. 369.

Best, Bill P.

Date: Mon Jun 22, 1992 9:20 am PST
Subject: Social chimps

[Bruce Nevin (920622) --

>I do not claim that Sagan and Drayan have correctly described the
>perceptions that chimpanzees are controlling.
...
>I do claim, and I think that you would agree, that projection of this
>kind is typical of what people do.
...
>However, instead of attempting to eliminate this process as an unwanted
>interference with "objectivity," I would embrace it as being itself a
>crucial datum about our perceptions and our control of them
...
>In particular, I make the further claim that the association of
>particular manners of behavior (comprising behavioral outputs), >emotional states, and social roles (participation in mutually >recognized social relations) is learned and indeed taught as part of >how to be an adult member of one's society.

All of these things you say are true; I agree, and apologize for having taken off from your line of argument in a different direction. Perhaps from my earlier post of this morning you may see what my objection is. I'm not objecting to the description (in somewhat less wildly projective form); only to the implication that this is the only way a human society could be organized. I wouldn't be surprised if many aspects of human society are there by default, carried over from remote ancestors and accepted simply because that's the way it's always been done. Also, as you say, these social interactions may come about through nonverbal processes, through each organism's trying to control what happens to it. By projecting our own perceptions into such situations, we learn more about what we perceive and control in these nonverbal ways.

But if we look at a chimpanzee society as what emerges from interactions without benefit of symbolic reasoning, principles, or system concepts, and if we see parallels (however described and interpreted) with human societies, isn't this a sign that human beings aren't really taking advantage of their own higher capacities? That's how I see such parallels -- not as evidence of some inescapable animal heritage, but as evidence of immaturity, of lack of skill at using brains in the ways they are now capable of being used. I don't blame a chimp for indulging in all that pomp and stuff; but when I hear of human beings doing the same thing, I wonder why they have to act like chimps.

This immaturity always leaps out at me from descriptions like those of Sagan and Drayan; for me to visualize human beings acting in these ways is to see a flawed and rather ridiculous mode of interaction based on illusion, false hopes, and misunderstanding of human nature. I can't imagine myself in any of the roles described. True, some people do behave in this way. But I am, I'm afraid, more interested in doing something about that than in studying it disinterestedly as a phenomenon. I have too much interest in it just to let it be. I can't just say, "Well, that's how society works and I guess we're stuck with it." I don't think we're stuck with society as it is. Or I don't WANT to be stuck with it. I REFUSE to be stuck with it.

Maybe this is why I never got a PhD. I just can't stick to the point.

Best, Bill P.

Date: Mon Jun 22, 1992 9:56 am PST
From: Dag Forssell / MCI ID: 474-2580
Subject: Leadership

[From Dag Forssell (920622-1)] Bill Powers (920622.0800)

>But leadership, it seems to me, is a role in a social hierarchy. It
>requires followers.

The idea that a leader is defined as someone who has followers is indeed the predominant interpretation in our society. I believe it an unfortunate one. It is not the only one available.

Webster's says:

LEADERSHIP n. 1. the position or guidance of a leader.
2. the ability to lead.

>I'm first of all an explorer of a new idea, second a teacher, and third
>a student. These are things that anyone can be. On this net, everyone
>plays these roles at one time or another. For one person to be any of
>these things does not rule out another person being the same things,
>at the same time, in the same group.

Yes, indeed. I think it fits nicely in the second definition above.

>The worst result, from my point of view and in my circumstances, is
>that followers learn from a leader how to follow, not how to explore,
>teach, and learn.

This does not follow from an emphasis on explore, teach and learn (a good model of reality), (applicable to any field of endeavor) which I think is a better definition of leadership.

>This new idea is far more important than I am; it will be remembered
>long after I am forgotten.

Yes.

>I don't want to think that when I disappear from the scene, my followers
>will cast around for someone else to follow, someone with another
>admirable idea, and will never think of carrying on for themselves what
>we are doing together now.

Again, this concern falls if you let go of the emphasis on following. I can't conceive of a control system wanting to follow. What a control system wants is good systems concepts to inspire good principles, so you can select effective programs ... so you can maintain your body chemistry. A control system is DESIGNED TO LEAD ITSELF; to satisfy its own purposes as it perceives them. Purposeful Leadership as I define it is the development and communication of good information that allows every individual to lead him/her self in full autonomy. It is a non-manipulative, non-coercive, non-violent approach. Your example from the newspaper was a pretty good example of this.

With good information shared and internalized voluntarily, people will be aligned and will automatically cooperate on the mutual concerns.

>I need friends, people I love and who feel love for me, colleagues,
>equals.

How about fellow leaders?

Not your blind "follower," but an autonomous student.

Dag

Date: Mon Jun 22, 1992 9:56 am PST
Subject: Leaders and followers

[From Rick Marken (020622.1030)]

Bill Powers (920622.0800)

>Thanks for the vote of confidence, but I'll stick to my definitions. I'm
>first of all an explorer of a new idea, second a teacher, and third a
>student. These are things that anyone can be.

> The worst result [of treating a person as a leader], from
>my point of view and in my circumstances, is that followers learn from a
>leader how to follow, not how to explore, teach, and learn.

I agree completely. What I said was:

>If what Bill does is Purposeful LeadershipTM then, boy, am I for it!!
>I just hope it can be taught.

I meant that if what you do is called Purposeful LeadershipTM then what Galileo and Newton did would also be called Purposeful LeadershipTm. That's why I thought it might be hard to teach it. How do you teach someone to be inquisitive and open minded and have a brilliant new way of understanding and exploring a fundemental aspect of our human experience?

>I don't want to think that when I disappear from the scene, my followers
>will cast around for someone else to follow, someone with another admirable
>idea, and will never think of carrying on for themselves what we are doing
>together now.

Fear not. You're awfully good, but the ideas are even better.

>I don't need followers. I need friends, people I love and who feel love for
>me, colleagues, equals.

You've got 'em!

I would like to see you recognized for the enormity of your insight (that behavior is the control of perception) and of your theoretical and scientific contributions. But that's for my sake (to control my perception of intellectual fairness), not necessarily for yours. And that's RECOGNIZED, not CANONIZED.

Your student and colleague

Rick

Date: Mon Jun 22, 1992 11:21 am PST
From: CHARLES W. TUCKER
EMS: INTERNET / MCI ID: 376-5414
MBX: N050024@univscvm.csd.scarolina.edu

TO: * Dag Forssell / MCI ID: 474-2580
Subject: MCI addresses

Dear Dag,

Someone gave me a MCIMAIL address that I think is wrong. Is there anyway other that contacting the person another way that I can get his address?

Thanks for you help.

Chuck

Charles W. Tucker (Chuck)
Department of Sociology
University of South Carolina
Columbia SC 29208
O (803) 777-3123 or 777-6730
H (803) 254-0136 or 237-9210
BITNET: N050024 AT UNIVSCVM

Date: Mon Jun 22, 1992 11:24 am PST
Subject: chimps, champs, chumps, etc.

[From: Bruce Nevin (Mon 920422 13:18:43)]

>if we look at a chimpanzee society as what emerges from interactions
>without benefit of symbolic reasoning, principles, or system concepts, and
>if we see parallels (however described and interpreted) with human
>societies, isn't this a sign that human beings aren't really taking
>advantage of their own higher capacities?

I don't know that chimps lack symbolic reasoning, principles, or even system concepts (or control of perceptions of the kinds we label with those terms).

>True, some people do behave
>in this way. But I am, I'm afraid, more interested in doing something about
>that than in studying it disinterestedly as a phenomenon. I have too much
>interest in it just to let it be. I can't just say, "Well, that's how
>society works and I guess we're stuck with it." I don't think we're stuck
>with society as it is. Or I don't WANT to be stuck with it. I REFUSE to be
>stuck with it.

Certainly I sympathize with this, a lot. But if your aim is to drive to Durango you have to first know if you're starting out from Chicago, or LA, or Boston, or Sydney.

But that metaphor is inept. I think it is inappropriate to imagine leaving behind all attributes of our primate, mammalian, reptilian, and other evolutionary forebears. Evolution doesn't seem to work that way. The lungs evolved from the swim bladders of fishes, and to their double membrane construction (wonderful for swim bladders) we may attribute the pleurisy from which my wife suffered last fall. Our throats represent a compromise between the needs of eating, the needs of breathing, and the needs of phonation. The horse, for example, optimized for breathing. Rube Goldberg make-dos are the norm rather than the exception. But unlike other critters (or at least more than they), we can consciously participate in our own evolution. However, even that concerns what we do with our inherited materials, and we ignore those materials at our peril.

So there is a goal of ameliorating our social relations with others, and not wanting certain chimp/chump attributes to be there. This conflicts with the goal of observing what is there dispassionately, without attachment to outcomes, just for the sake of finding out whether we're starting out so to speak from New York City or Dry Prong, Wyoming, a prerequisite for the goal of amelioration. Perhaps it is advisable to set aside the emotional craving that those undesired attributes not be there, at least at the time of taking observations.

What might the positive useful functions of such chimpish shenanigans be? Look at your account of the elevation of the clerk to the manager's slot. You accorded him some of the outward signs of status, such as calling him "Boss." These signs may have "felt good" to him, and it may even on occasion have "felt good" for one of you to "make" him "feel good" in this way. More significant, I think, from the perspective of an outsider must have been the perception of one person answerable for your department. Even if it is known that this person would just go and get the answers from one or more of you, he knows whom to get it from and how to go about getting it in a form that is useful to the particular outsider, who doesn't want to know any great detail about all the stuff you guys are maintaining. He knows (or can learn) what the expectations and needs are of upper management, other departments, outside vendors, competitors, reporters doing a story, directors, and so on, a proliferation of PR requisites from which you guys in turn want him to insulate you. I'm just imagining and projecting, of course, and I may well be way off target on some particulars, but the general idea I think comes across?

But this is simply specialization of function, you might say, according to skills, training, temperament, and other things that would lead one (say, in an anarchist society) to prefer helping one's fellows in one way rather than in another. There is no hierarchy inherent in it. Well, no, but there is. It is a hierarchy that follows from something like "chunking" of information. But I agree that there is no dominance or coercion, aka "power," inherent in it.

An individual contributing in a hierarchically superordinate position (in this "chunking" sense) is able to take unfair advantage of that position more effectively than one in a less centralized position, on which fewer others are so dependent. Then a question becomes, for what reasons might such an individual come to abuse the strategic advantages of the position? Your "boss" evidently realized his dependence on those who had elevated him to the flak-catcher position, was grateful for it, and was conscious that he could not maintain himself there without your support. Conversely, if he abused his position and seemed irremediably abusive, the rest of you could withdraw your support in many ways. This is what Gene Sharpe's many years of research on nonviolent resistance brings out, the many forms this withdrawal of support can take, and how surprisingly effective it can be against the most tyrranical rulers.

The point of this little disquisition is that it does not seem to be so simple as atavistic hormones inducing us to chimpanzee behavior, this business of people seeking "leadership." Part of it is the leader as representative, including the various flak-catcher roles. In the extreme, the leader becomes the representative sacrifice. And that is another reason why some of us have turned down the role and its benefits.

Another dimension not to be overlooked is the hierarchical power relation between adults and small children. No amount of egalitarian system concepts, principles, and programs can undo the biological facts of dependence and weakness and need for nurturance. Bateson's wolf pack example is telling here. It is the mother's gesture while weaning her puppy that the pack leader makes to the usurper of sexual access. This communication is powerful and effective because of the direct analogy it makes between the leader-member relationship and the mother-puppy relationship, and between the usurper's present situation and a critical juncture in the development of the latter.

What is analogy?

Bruce bn@bbn.com

Date: Mon Jun 22, 1992 11:38 am PST
Subject: follow the leader

[From: Bruce Nevin (Mon 920422 15:02:32)]

(Dag Forssell (920622-1)) --

>I can't conceive of a control system wanting to follow.

Oh, come on, Dag! You can't mean that, can you? Aren't there many occasions when one control system wants to follow the lead of another control system? And is this in itself pernicious? (Though it can be abused--on both sides of the dyad, be it said! Nor does it end with childhood. Nor is it always childlike, though abuse of childrens' dependency does seem to result in many adolescents and adults comng to abhor and scorn it and fear exposure of it in themselves. One of the sure recipes for childishness.)

Have you ever taken a dance class?

Bruce bn@bbn.com

Date: Mon Jun 22, 1992 1:59 pm PST
From: marken
EMS: INTERNET / MCI ID: 376-5414
MBX: marken@aero.org

TO: * Dag Forssell / MCI ID: 474-2580
Subject: letter

Dag

I think the letter is really looking good. It starts really well.
I haven't gone over it in detail but it scans well -- it should get some bites.

Best regards Rick

Date: Mon Jun 22, 1992 3:50 pm PST
Subject: Re: Sameness; chimps

[Martin Taylor 920622]
(Bill Powers 920618 ?)

Yesterday was the first day of summer. Temperature maximum in Toronto 10.8C (about 51F), breaking the previous record for coldest maximum by 4 degress C. Has Mount Pinatubo overcome the global warming?

Bill addresses the problem of tracking multiple "same" objects, but I think he still sidesteps the issue Bruce brought up initially, that I was trying to buttress with my introduction of the tracking study. First, on the tracking:

>The problem you state breaks down into two questions: first, how does a
>person keep track of SINGLE individual in a set of identical moving
>objects, and second, given that the first question is answered, how can
>more than one such process occur in the brain at the same time?

Bill points out that in principle one can only the objects if the targets maintain differences from the distractors in at least one of position, velocity and acceleration. That's quite true, and experimentally verified. But even after two have been confused, the observer can still track about four (depending, I think, as Bill says, on the motion statistics of the ensemble). One of the four may be, from the experimenter's viewpoint, wrong, but the observer sees the one that best fits whatever tracking method is used by the ECSs that are employed. To be brief, I think Bill's analysis of these two questions is more or less correct. But...

As I see it, the key problem is more like: third, how does the brain identify that there exists N objects of a certain kind rather than one strong exemplar of the kind. I've forgotten how Bruce originally put it, but perhaps I can paraphrase. If there is an ECS controlling for the perception of an X, it should satisfy its reference if an X is in its input. What distinguishes thes existence of exactly three Xs in its input from the existence of one X. Bruce asked if there might be a multiplicity of ECSs controlling for an X, which would solve the problem if each was controlling for "an X THERE" rather than just for an X. If THERE were some kind of an internally generated reference (like a memory--a novel construct at this point, which I am not going to stand behind), then the discovery of an X could lock THERE to the place where the X happened to be, and would permit tracking. But how could any other ECS controlling for the perception of an X know that the first one was "taken?" I think there is a real problem here, and none of Bill's proposals
to date seem to solve it.

Please note that I am neither expecting nor demanding that Bill solve every problem of PCT. Noting his comments on leadership, I hope that other people (perhaps including myself) can solve them. But problems such as this one are fairly central, I think, and must be solved within the "natural" hierarchic structure if HPCT is to be taken seriously as the instantiation of PCT that corresponds to real living beings. The basic statement of PCT, that behaviour is all and only the control of perception, seems incontrovertible. How that fact is developed into structure is not. HPCT seems a very sensible proposal, but there may well be other equally sensible instantiations, and not all living things necessarily use the same instantiation.

We accept as a working hypothesis that the perceptual control structure is layered. Bill said in some posting shortly after I joined this group that whenever he had thought of a level-jumping control, it turned out not to be (appropriate/correct/necessary/simple?). Maybe so. In the neural network business, multilayer perceptrons often work well, but there are other architectures that are more appropriate for complex problems, including specifically gated modular architectures in which smallish modules solve sub-problems, and gating structures determine which subnetworks present their solutions to higher modules. Perhaps control nets might work better on complex
problems if conflicts can be resolved by gating structures that permits some modules but not others to exercise control? There are multitudes of possible architectures, and we may not be constructed to use the most obvious one.

Simplicity is to be preserved where possible in science, and HPCT is a simple structure. If it can solve problems as basic to perception as "there's an X and there's another" or "I see a lot of Xs among the Ys" without the introduction of new structures or concepts, so much the better. At present I don't see the answer. Martin

Date: Mon Jun 22, 1992 4:01 pm PST
Subject: Re: looking at looking at chimps

[Martin Taylor 920622 19:30]
(Bruce Nevin and Bill Powers various postings on Sagan and society)

I have thought that one of the great insights from PCT was that the experimenter is a control system that could control those external variables that might be being controlled by the subject of the experiment. Only if the experimenter can perceive the "same" environmental complex as the subject does can the experimenter apply the TEST. This REQUIRES anthropomorphising the subject. One cannot perturb the variable that might be controlled by the subject unless one can perceive it AND control it oneself.

In an observational science, the experimenter cannot perform the TEST in the real world, but can do it in imagination, can perceive disturbances that the environment applies to the possibly controlled variable, and see whether the observed subject acts so as to return the variable to some (presumed) reference state. The experimenter has to take the place of the subject in imagination. Anthropomorphizing is not only legitimate, but necessary. If you don't do it, you are no better than an S-R psychologist.

In that light, the Sagan-Drayan description is a much better description of the "royal chimpanzee" than would be a description that the alpha male drew back his lips to expose his teeth and raised his hand, and that another male lowered his head immediately afterward. Doesn't that kind of a description sound like the "behavioural" descriptions against which the PCT "leaders" so often rail? (Does a loose canon run on rails? I think not, as a PCT mind cannot be one-track).

Martin

PS.
Bill

>I've ranted for years to the CSG that if we want to have a revolution, we
>must revolt.

I don't find CSG revolting, and I don't want to. Jolting is OK, but revolution seldom has good results, politically or scientifically. There's lots of good food for thought out there among the garbage. Some just needs to be made a bit more tasty by being taken with a grain of salt.

Date: Mon Jun 22, 1992 4:36 pm PST
From: Dag Forssell / MCI ID: 474-2580
TO: Chuck (Ems)
EMS: INTERNET / MCI ID: 376-5414
MBX: N050024@univscvm.csd.scarolina.edu
Subject: MCImail
Message-Id: 33920623003633/0004742580NA2EM

Your message gave me your internet address as:

N050024@univscvm.csd.scarolina.edu

So if this works, you can add it to your address file.

MCImail has a customer service phone number: 800-444-6245. They will look up a number for you if you have an exact name. You can perhaps discuss your incorrect # with them. They will of course look up on a screen. Note that you would not find Greg Williams that way. He is listed as Hortideas. So be prep- ared to be imaginative.

Good luck. Dag

Date: Mon Jun 22, 1992 4:49 pm PST
Subject: Leaders, followers

[From Bill Powers (920622.1600)]

Bruce Nevin (920622.1318) --

> I don't know that chimps lack symbolic reasoning, principles, or even
>system concepts (or control of perceptions of the kinds we label with
>those terms).

Well, neither do I. It would surprise me, though, to find chimps capable of the same degree of skill at these levels that are potential in humam beings, especially considering the limitations on communication of ideas in a COMPARATIVELY languageless culture.

>But if your aim is to drive to Durango you have to first know if you're
>starting out from Chicago, or LA, or Boston, or Sydney.

Good point. On the other hand, having decided to go to Durango, I don't need to do a poll to see if that's where everybody goes. Nor do I need to devote a lot of time studying the routes to places where I don't want to go and that I don't recommend to others.

>I think it is inappropriate to imagine leaving behind all attributes of
>your primate, mammalian, reptilian, and other evolutionary forebears.

It is, I agree, important to understand the means of transportation. However, I don't believe that we are primates riding on horses riding on reptiles. All levels of organization have evolved on our way to being human, not just the highest. As human beings, we are not limited to the reptile's means of thermoregulation or child care, for example. We are more skillful at most motor control processes than are horses or even monkeys and chimpanzees. A man can, eventually, run down a horse.

>The horse, for example, optimized for breathing.

Nonsense. A horse is optimized for betting.

>So there is a goal of ameliorating our social relations with others,
>and not wanting certain chimp/chump attributes to be there.

Social relationships that are properly understood from the outset don't need ameliorating. But I agree with your basic thesis here, which is that we have to understand where here is before we can get to there. And I also agree that there is value in some aspects of functional organization, especially when many people are trying to accomplish something complex together.

The point I was addressing was the confusion of SOCIAL position in an organization with FUNCTIONAL position. V. was in a position functionally, hierarchically, superior to that of everyone else in the department. We all agreed that that function needed to be carried out. For the good of the paper, which we considered to be our own good (working at a newspaper is rather neat) we agreed to act, in certain matters, as if V. were a higher-order control system. If he decided that too free a swapping of vacation time didn't fit in with the newspaper's general personnel policies, as he did at one point, we grumbled but conformed, because he had to deal with other departments and we didn't. In other policy matters we simply said that we would not work under the proposed policy, and worked out a compromise. All these relationships were functional and unrelated to social position.

>You accorded him some of the outward signs of status, such as calling him
>"Boss." These signs may have "felt good" to him, and it may even on
>occasion have "felt good" for one of you to "make" him "feel good" in
>this way.

Yes. V. valued such things not because they meant anything inside the department but because they meant a lot in his dealings with managers outside the department, where social ordering was pronounced. We picked him in part because he WAS concerned with social status and smarted under the burden of being ignorant in a subculture that valued competence. We solved his problem and ours at the same time, and it turned out to be a rather elegant solution -- functionally from our point of view, and socially from his.

>But this is simply specialization of function, you might say, according to >skills, training, temperament, and other things that would lead one (say, in an >anarchist society) to prefer helping one's fellows in one way rather than in >another. There is no hierarchy inherent in it. Well, no, but there is. It is >a hierarchy that follows from something like "chunking" of information. But I >agree that there is no dominance or coercion, aka "power," inherent in it.

You got it. That's my design for a better social system.

>An individual contributing in a hierarchically superordinate position
>(in this "chunking" sense) is able to take unfair advantage of that
>position more effectively than one in a less centralized position, on
>which fewer others are so dependent.

Sure, if the others let it happen. But once people learn the distinction and come to see the advantages of the functional approach, their resistance to that kind of disturbance is going to be pretty implacable. The greater the numbers of people affected, the quicker a little resistance by each one adds up to a brick wall.

>Another dimension not to be overlooked is the hierarchical power
>relation between adults and small children.

Ah, yes. This is why I likened the relationship between leaders and followers to one between adult and child. This doesn't have to be carried out as a power relation most of the time -- only when the child persists in incurring danger for self or others. It is possible to respect the will of the child and to avoid making an issue of relationships that have value ONLY in a power structure (like "don't talk back to me").

Also, keeping Eric Berne in mind, there's nothing structurally unsound in deliberately adopting the role of child, and saying to another, "Tell me what to do." If I were to try to learn to water-ski, that's about how I'd have to begin. When I'm discouraged, I want to be comforted. It's the system concept under which all these detailed actions and relationships take place that makes all the difference. The kind of childhood one had makes all the difference in how one takes refuge now and then in the role of child.

John Gabriel (off-net, copy to CSGnet)

>I really like your posting. The fifteen happiest years of my lfe were
>spent in a research group run very like the newspaper group you
>describe. But perhaps on reflection you might want to change some
>terminology or emphasis. Conflict always arises around new ideas, and
>Darwinian selection whether of ideas or species really is red in tooth and claw.

The conflict is worst when the people with different ideas are competing for social position. When being right is confused with being valuable, differences in theories reflect on one's self-concept. If you know that you're valued and considered an equal, the flavor of intellectual debate changes completely. It's only when people use intellectual interactions as means for gaining control or power over others, when they seize on others' mistakes as a way of clawing their way past them on the social ladder, that the nastiness begins and the joy departs.

When social competition is removed, it becomes safe for the author of a theory to become his/her own most severe critic. Admitting a mistake in public is then no big deal; it doesn't indicate any unusual dose of "character." It's just what you do when you make a mistake, if someone else hasn't found it first. When you don't have to be defensive of your own ideas, you're the most likely person to see what is wrong with them, because you've probably thought about them more than other people have and are the most likely person to catch subtle errors. In the right kind of atmosphere, you can take a sort of Zen-like attitude toward being right -- it matters, in a way, but it doesn't really MATTER.

It's hard to maintain this attitude, of course, when you're surrounded by people taking notes, who are just waiting for you to make some silly mistake so they can jump on you. The only way I know of to avoid such treatment is first to avoid giving it to others, and second to choose your friends and associates wisely.

(Dag Forssell (920622-1)--

I'm still a little leery of selling "leadership," not because I object to your understanding of it but because your understanding may not be the one at the receiving end of your letters. As Bruce pointed out, and I echoed, all of us occasionally, and willingly, follow. But that's a temporary condition, not an essential part of the system concept. Now that the subject has been mentioned, I can think of many times when one person or another at the paper, who knew the nature of the problem the best, fell automatically into the role of leader: "We need another memory board, and make sure it's been tested. And we'd better get another disk drive ready as long as the system is down." And everybody else snaps to it. The next time it will be someone else calling the shots. In fact I can recall times when too many people pitched in enthusiastically and produced chaos; then someone would say "Joe, you take over, we're getting in our own way."

Your new letter has places in it that give that leap of the heart that says there's something real here. I hope some of the recipients get it.

I'd delete adjectives wherever possible: contrary to beliefs in the advertising industry, saying "a bright new future" is weaker than saying "a new future." Think of the impression left when you answer the question, "Is this person able to do this job?" with "Yes, he's really very competent," as opposed to "Yes. He's competent." Or even just "Yes."

Best to all, Bill P.

Date: Mon Jun 22, 1992 7:03 pm PST
Subject: Re: AI, PCT, and HPCT

I haven't been keeping up with CSGNet this summer, since I am "off" (read "unemployed") from teaching until August, so I have to dial in to the net from home, which is expensive (long-distance). I check my mailbox only once a week or so, so forgive my late remarks on AI.

I realized some time last Spring that the HPCT way of looking at things was coming to dominate my conceptualization of human cognitive activity. Since I teach and do research in AI/Cognitive Science, this implies that I am in the middle of a paradigm shift, and am having to re-cast my theoretical foundations for my research in HPCT terms. Trying to do this on the fly has become confusing to me, so I am in the process of jotting down some loosely structure notes that are intended to remind me of how things work, and what kinds of things I ought to be l looking for, from within the HPCT paradigm. From what I have observed so far, I think it is clearly possible, and fruitful, to do AI from the HPCT perspective.

However, there are some difficulties. One of my problems is to account for the source of the reference signals; are they innate, or are they acquired by the organism through interaction with its environment? If the latter, some other reference signal must have been present prior to the acquisition of THIS reference signal, because all perception is mediated, a product of external stimuli and internal reference signals. Fortunately, it seems to be the case for color classification and naming that there IS an innate set of reference signals that come wired in to the human organism at birth. This suggests that, although we later have to pull ourselves up with our bootstraps, at least we are born wearing boots with straps. (I believe Kant suggested this a while back.)

There are two types of AI researchers: those who want to make machines "intelligent" so they can perform some task better, and those who want to make machines perform some task like humans so they can] better understand what "intelligence" is. I guess I'm more the second type. If I can design an AI system that utilizes an HPCT approach to model some intelligent (human) behavior, and it works better than systems that don't use an HPCT approach, then this may tell us more about how human intelligence works. Right now, Bill Powers (and the other folks who are doing simulations) is really doing AI research (whether he likes it or not!) Currently, robot arms aren't very "intelligent" - for various hardware reasons, but also because they are based on a faulty model of the human arm. Bill's conception of how the human arm works is "better" (in the sense that the model responds more like human arms do), so an actual physical robot arm built according to his simulation should work better than current robot arms. This is true AI (of the first type).

Unfortunately, not much AI is actually based on ANY model of human intelligence; "whatever works" seems to be the watchword. So I don't know that HPCT will have an overwhelming impact on the field; but it could, and it should. It is certainly affecting how I think about my own AI work.

- Gene Boggess

Date: Mon Jun 22, 1992 4:36 pm PST
From: Dag Forssell / MCI ID: 474-2580
TO: Chuck (Ems)
Subject: MCImail

Your message gave me your internet address as:

N050024@univscvm.csd.scarolina.edu

So if this works, you can add it to your address file.

MCImail has a customer service phone number: 800-444-6245. They will look up a number for you if you have an exact name. You can perhaps discuss your incorrect # with them. They will of course look up on a screen. Note that you would not find Greg Williams that way. He is listed as Hortideas. So be prep-ared to be imaginative.

Good luck. Dag

Date: Mon Jun 22, 1992 7:46 pm PST
From: Dag Forssell / MCI ID: 474-2580
Subject: Leadership

[From Dag Forssell (920622-2)]

Bruce Nevin (920622.1318), Bill Powers (920622.1600)

Ok, Bruce. A difficulty on this net is that anything can be and is taken so damn literally. You have a point of course. I did learn to lead in dance once upon a time. It is important in ballroom dancing to give clear signals to your partner (follower) with a steady hand. The follower chooses to follow and concentrates on that.

Leadership is often understood to mean that you tell someone WHAT TO DO, then they follow by doing what you tell them to do. This emphasis on doing and instructions fits nicely in a cause-effect world.

I am trying to redefine and sell Leadership as the idea that if you want to lead, the MOST EFFECTIVE way is to offer good (a description of that Boss Reality that is as good as you can make it or negotiate it in open discussion) information for your "followers" to evaluate and make part of their own systems concepts if they want to (understand, no conflicts with pre-existing concepts, relevant etc). --- Then you step out of the way and let the "followers" control to their hearts content. You will not need to supervise or "control" their actions, because that is built in.

This is what I read into that quote from Max De Pree.

This form of leadership is inherently non-violent. Teaching it will not work if the top management in a company is coercive, as I perceive most to be. Therefore the idea MUST be sold at the very top. To the very people who are used to insisting on results or else. I am counting on finding a few who will see it my way, but don't expect many. A few is all I need. Once the process is understood, the leadership/information can come from anyone in the organization.

This net offers leadership from all corners, causing re-organization and growth and progressively better understanding.

I appreciate Bill's kind comments on the letter. I shall review it one more time for adjectives. Christine has already been after me for months on that, and I thought it was fairly free from unsupportable judgements.

Dag

Date: Tue Jun 23, 1992 7:28 am PST
From: CHARLES W. TUCKER
EMS: INTERNET / MCI ID: 376-5414
MBX: N050024@univscvm.csd.scarolina.edu

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Re: MCImail

Thanks Dag for the info; I will try it.

Regards, Chuck

Date: Tue Jun 23, 1992 9:38 am PST
Subject: fish food for thought

[From: Bruce Nevin (Tue 920423 12:34:48)]

Over lunch I just read a summary of research on social hierarchy and testosterone in an African species of fish (_Richard Francis et al. in _Ethology_ 90:247, summarized in _New Scientist_ for 6 June 1992 p. 15).

Dominant males bear certain colorful markings, have larger testes and higher levels of androgens like testosterone, and can mate; other males lack the markings, have smaller testes and androgen levels, do not mate, and "the entire system that links the brain to the pituitary gland to the gonad" is suppressed. The territorial males "aggressively dominate" the other type "even if they are of similar size."

When a territory-controlling male dies, "a drab male will quickly become colorful and behave just as nastily as the previous owner."

The question asked of an experiment was, which comes first: aggressiveness, social superiority, or large testes.

They castrated dominant males. As expected, levels of testosterone went down dramatically, and the fish were less aggressive. However, "in most cases, the castrated and less aggressive males retained their dominance over intact, non-territorial males, even after they had been separated from each other for six weeks."

The researchers draw the inference that "prior experience plays an important role in determining dominance." The conclusion of the NS summary says in part:

> Apparently, being dominant makes you more likely to be dominant in the future. The experience can alter parts of the brain which in turn control the production of . . . gonadotropin-releasing hormones, or GnRH . . . which then regulate gonadal (testicular) development.

It would seem that the biological regulation of social structure among autonomous control systems is pretty basic and pervasive.

I wonder what would happen if they painted a non-dominant fish to look dominant. Would it be attacked? Might its response include becoming in fact one of the dominant type?

Would probably have to remove a dominant male (as if in the case of one dying). Would some part of the ensemble of markings suffice for the other fish to treat the one with makeup as incipient successor to that territory, and would that in turn result in our nominee getting the slot? Or would such human favoritism be overridden by an emergence of greater aggression on the part of one of those perceiving the vacancy? Perhaps it takes a degree of readiness; it seems likely that there is an unnoticed intermediate level in the hierarchy, of fish ready to pop into a vacated slot.

Hormones, neuropeptides, etc. seem to interact with neural control systems (and perhaps vice versa) by way of receptor sites in close proximity to neural structures. They may have much to do with readiness or even ability to notice things like a vacancy in fish territories.

Bruce bn@bbn.com

Date: Tue Jun 23, 1992 10:22 am PST

Subject: different cat

[From: Bruce Nevin (Tue 920423 08:32:46)]

(Martin Taylor 920622) --

>As I see it, the key problem is more like: third, how does the brain identify
>that there exists N objects of a certain kind rather than one strong exemplar
>of the kind.

One possibly fruitful way to approach this is by attributing differentiating perceptions to the individuals in memory and imagination. This is what we do with language ("the green one"), and is kind of an inverse of how Penni's "say-elaborating" process works. What I am proposing is that we do this prior to language (though we use language to facilitate and elaborate it). Indeed, I have proposed that it is out of a clustering of attributed perceptions that categorization itself arises.

Suppose we have ECSs (elementary control systems) controlling for ways to differentiate between like perceptions by identifying other perceptions associated with them as attributes. Is there any reason this is implausible? In such a way we might develop control of category perceptions.

One way to track four dots that soon becomes intractible is using their history ("the one that started in the northwest corner"). An ECS alert for differentia might entertain various hypotheses ("the bounciest one") and develop a history of them (". . . is tired now"). A story, not in words but in remembered perceptions. A story often told becomes a stereotype. A stereotype of remembered behavioral outputs, perception of any part of which may result in the expectation of the remainder (or the sequel).

I think one way we differentiate between likes is subsumed into another process, one that is more general because it is not limited to like perceptions, one whereby perceptions are construed as elements in a structure or system. For example, if the movements of four dots among others were not entirely random but were such that they always formed some sort of quadrilateral with the same individual in the northwest quadrant, etc., the tracking problem would be much easier. If from time to time the higher level structure (in this simple case a configuration perception) were to change, tracking by this means could continue by controlling for the history of change at the higher level. One might detect and control for a cycle or rhythm in these changes (I believe we must have some inbuilt predisposition to control for cyclicity).

As we have discussed in the past, something that is structurally superordinate is not necessarily superordinate in the perceptual hierarchy. Two configuration perceptions going on concurrently, in parallel, themselves each provide input to a third configuration ECS. The interaction of the two sorts of hierarchy gives the appearance of loops on the same perceptual level (we have configurations of configurations, idioms are sequences of words, which are sequences). There may be loops involving more than one level of the control hierarchy. I don't want to open that can of worms now (we're so easily distracted), just to put out the hunch that this proposal might provide means for dealing with these coherently.

I believe that creatures evidently without language, such as chimpanzees, possess these capacities. Language gives us not so much a new capacity, as finer control of it because of (among other things) longer, easier, and clearer retention in memory, elaboration of these capacities to higher levels of the control hierarchy, and above all transmissibility. It is because of the transmissibility of linguistic information that we can be less constrained to re-enacting stereotyped behavioral outputs for the sake of social cooperation, as compared with primates and other critters.

Which broaches the question of the basis for trust. Trust is the foundation of leadership in the functional sense and authority in its non-coercive senses, and trust is the reciprocal of coercion, that is, lack of trust is the occasion and "justification" for leadership and authority in their coercive senses. The evidence that I have seen suggests that trust depends strongly on those mammalian and more primitive stereotypes and expectations being fulfilled, including things studied as body language and gestural manners-of-use of ordinary language (conversational style per Tannen, etc.). These are reliable precisely to the extent that they are *not* subject to conscious manipulation. Nor is there any especial need to "outgrow" them per se: they are not to blame for abuses of social assymmetry by individuals whose trust as children was abused by adults who were abused, ex antiquo, and other ripples into our biological present from various ancient social dislocations, ancestral experiences of expulsion from the Garden where we all, did we but awaken to it, in fact still dwell. Among the memories and imaginings by which we differentiate individuals, it is those bad dreams that we need to outgrow. For that, trust is both the key for opening and the key that is found upon opening--rather like the nested "aha!"s of science, or of art.

And that is one reason we do so well here. So how does the HPCT perspective engender trust? To say you can't do otherwise than let other control systems control because they will do so anyway is not exactly reassuring to those who live in fear and therefore want to control others. I have some thoughts here, but I'll let them rest.

Date: Tue Jun 23, 1992 11:12 am PST
Subject: Re: Sameness

[From Rick Marken 920623.0945)]

Martin Taylor (920622) says:

>As I see it, the key problem is more like: third, how does the brain identify
>that there exists N objects of a certain kind rather than one strong exemplar
>of the kind.

Very interesting question.

Let me make sure I've got it. Is the question something like " How
can I recognize a repetition of the same pattern - like AAAAA? Doesn't
the "A" detector get "used up" with just one?"

One obvious answer is that there must be several detectors for the
same sensation, configuration, transition, etc. How this is architected
is the big question. How does a model like "pandemomnium" solve this
problem? Or does it?

>If there is an ECS controlling for the perception of an X, it
>should satisfy its reference if an X is in its input. What distinguishes
>thes existence of exactly three Xs in its input from the existence of one X.

Again. I'd guess that "three X's" requires three X detectors firing
maximally and simultaneously. The "threeness" must be detected by a
count detector that takes the X detector outputs as inputs. How that's
done, I don't know. Looks like a research/modelling problem to me.
Get those grad students on it, Martin.

> But problems such as this one are
>fairly central, I think, and must be solved within the "natural" hierarchic
>structure if HPCT is to be taken seriously as the instantiation of PCT that
>corresponds to real living beings.

How are these "central" problems solved by existing cognitive models? These
models might give some hints about possible mechanisms. But I think detailed
research might be the best first step. For example, how about a study
in which a person controls N of the same objects simultaneously? See how
control varies as a function of N. The problem would be to keep the number
of output degrees of freedom constant while the perceptual degrees of
freedom increase.

>   The basic statement of PCT, that behaviour
>is all and only the control of perception, seems incontrovertible. How that
>fact is developed into structure is not.

Right! But the idea of behavior as controlled perception is THE important
concept -- the resulting shape of the PCT structure will start to emerge
once there are more than three people in the world who are doing research
and modeling based on an understanding of the fact that behavior IS control.

>HPCT seems a very sensible proposal,
>but there may well be other equally sensible instantiations, and not all
>living things necessarily use the same instantiation.

No doubt.

>Simplicity is to be preserved where possible in science, and HPCT is a simple
>structure. If it can solve problems as basic to perception as "there's an
>X and there's another" or "I see a lot of Xs among the Ys" without the
>introduction of new structures or concepts, so much the better. At present
>I don't see the answer.

There are TONS of unanswered questions -- fortunately for those of use who
don't have billions of dollars to entertain ourselves with and, thus, are

forced to get our kicks from by trying to understand nature. I don't see the immediate lack of a definitive answer to the question "how do you see a lot of Xs among the Ys" as posing any more fundemental challenge to PCT than it does to S-R or cognitive models. It's just another interesting question -- and one that is far more interesting when asked in the context of an organizational model of behavior that actually works (PCT).

But I think there should be something we could learn from existing models, like neural networks. After all, such models are just models of pieces of a control organization -- posing as a complete model of behavior.

By the way, Martin (and other speech/language afficionados) -- any comments about the sound adaptation study that I mentioned in an earlier post?

Regards Rick

Date: Tue Jun 23, 1992 11:38 am PST
Subject: Re: Sameness

[Martin Taylor 920623 13:30] (Rick Marken 920623.0945)

>>As I see it, the key problem is more like: third, how does the brain identify
>>that there exists N objects of a certain kind rather than one strong exemplar
>>of the kind.
>
>Let me make sure I've got it. Is the question something like " How
>can I recognize a repetition of the same pattern - like AAAAA? Doesn't
>the "A" detector get "used up" with just one?"

I'm no longer trying to restate what I thought Bruce was getting at, but speaking for myself alone...No, that's not really the question. More like how do I know there are three A's in BAbNvXCA3kiAlaft without explicitly counting them in serial order, bringing each into central attention separately.

>How does a model like "pandemomnium" solve this problem? Or does it?

I know of no satisfactory neural net proposal that solves this problem.

>Again. I'd guess that "three X's" requires three X detectors firing
>maximally and simultaneously. The "threeness" must be detected by a
>count detector that takes the X detector outputs as inputs. How that's
>done, I don't know. Looks like a research/modelling problem to me.
>Get those grad students on it, Martin.

I don't think this solves the problem. How do these "X" detectors know that they are each looking at a different examplar?

I don't have students. I work in a government lab, and can engage contractors on research if the authorities agree that the problem is worth spending money on.

>> But problems such as this one are
>>fairly central, I think, and must be solved within the "natural" hierarchic
>>structure if HPCT is to be taken seriously as the instantiation of PCT that
>>corresponds to real living beings.
>
>How are these "central" problems solved by existing cognitive models? These
>models might give some hints about possible mechanisms. But I think detailed
>research might be the best first step. For example, how about a study
>in which a person controls N of the same objects simultaneously? See how
>control varies as a function of N. The problem would be to keep the number
>of output degrees of freedom constant while the perceptual degrees of
>freedom increase.

I'm not trying to compare HPCT with any other model. I take it to be a claim of HPCT that it is the "theory of everything" for psychology, and indeed, I have been trying to turn people on to it by saying that HPCT is to psychology what Newton's laws were to physics 300 years ago. What I am trying to do is to justify that claim, because I at present believe it to be true. Hence...
> I don't see
>the immediate lack of a definitive answer to the question "how do you
>see a lot of Xs among the Ys" as posing any more fundemental challenge to
>PCT than it does to S-R or cognitive models. It's just another interesting
>question -- and one that is far more interesting when asked in the context
>of an organizational model of behavior that actually works (PCT).

is perhaps true, but somewhat off the point. Seeing multiple things is an obvious and low-level ability that sufficiently complex living things have. It is at the level where HPCT should be able to provide a natural answer. Sure, there are tons of questions. It might be a good idea to develop an inventory of ones that seem simple to pose and that should be answerable either by analysis or through experiment within the (H)PCT paradigm.

Martin

Date: Tue Jun 23, 1992 12:32 pm PST
Subject: VOT experiment

[From: Bruce Nevin (Tue 920423 14:07:32)]

(Rick Marken (Tue, 23 Jun 1992 09:46:35 PDT)) --

>By the way, Martin (and other speech/language afficionados) -- any comments
>about the sound adaptation study that I mentioned in an earlier post?

You mean this one:

(Rick Marken (920621)) --

>The reviewer pointed me to a book that is, indeed, relevant to
>the topic of my paper -- it's called "The organization of
>perception and action" (1987) by D. G. MacKay (a UCLA linguist. . . .
>One interesting study that MacKay mentioned, which stongly
>points to the importance of feedback (though he didn't see it
>that way), showed that adapting a person with repeated presentation

>of the sounds /pi/ or /ti/ lead to a decrease in voice onset
>time (VOT) when these perceptually adapted people were asked to
>say /pi/ or /ti/. The study was done by Cooper and Naper and reported
>in JASA (Journal of the Acoustical Society of America) in 1975.
>Any linguists know about this? It sounds like there is a change
>in output (VOT) in order to produce an intended perceptual result
>(/pi/ or /ti/) via an adapted sensory input system. Is this
>interpretation reasonablee? It would have been nice if the
>study had been done quantitatively -- maybe it was. For example,
>the degree of change in VOT might be expected to depend on the
>degree of adaptation.

I didn't understand this. If "decrease" in VOT means (as I think it does) that voicing started sooner than in the corresponding unadapted syllables /pi/ and /ti/, this sounds like adaptation results in a relaxation of (reduction in gain for) control of contrast with the corresponding voiced sounds (less contrast between adapted /pi/ or /ti/ and ordinary /bi/ or /di/, resp.). Alternatively, maybe the voiced /bi,di/ would shift to relatively earlier VOT too--more "fully voiced"--preserving contrast with shifted boundary. You ask whether your interpretation is reasonable or not, but I am not clear what your interpretation is. How is such a change purposeful "in order to produce an intended perceptual result (/pi/ or /ti/) via an adapted sensory input system"? What am I missing about the nature of adaptation? Martin, can you help here?

I will look for the JASA article when I get the chance. Do you have access to it?

Bruce bn@bbn.com

Date: Tue Jun 23, 1992 1:18 pm PST
Subject: VOT experiment

[From Rick Marken (920623.1200)]

Bruce Nevin (Tue 920423 14:07:32)

>You mean this one:

>of the sounds /pi/ or /ti/ lead to a decrease in voice onset
>time (VOT) when these perceptually adapted people were asked to
>say /pi/ or /ti/.

Yep.

>I didn't understand this.

I had the same problem.

>You ask whether
>your interpretation is reasonable or not, but I am not clear what your
>interpretation is.

Just that there is an apparent change in output in order to keep a sensory result constant in the face of disturbance (in this case the disturbance being the adaptation). I just didn't see how a reduction in VOT could preserve the /pi/ or /ti/ perception so I was wondering if there was some esoteric acoustical phonetic reason. I guess I'll have to look at the original article to get a better description of what really happened. But whatever it was I can't imagine why some characteristic of speaking (VOT) would change after sensory adaptation unless it were being done to control a perception. So I was looking for some acoustic/phonetic evidence that would make the connection clear. Maybe the VOT change is just a response to the adapting stimulation? That would be a nice result for open loop fans.

I wonder how the authors of this work explained their results?

I'll try to look at the original JASA article today.

Regards Rick

Date: Tue Jun 23, 1992 3:02 pm PST
Subject: VOT adaptation question

[From: Bruce Nevin (Tue 920423 15:24:42)]

(Rick Marken (920623.1200)) --

>there is an apparent change in output in order to keep a sensory
>result constant in the face of disturbance (in this case the disturbance
>being the adaptation).

Is adaptation a disturbance? (Serious question. I don't understand what is going on in situations to which people apply this term.)

>I just didn't see how a reduction in VOT could
>preserve the /pi/ or /ti/ perception

Depends on how much onset of voicing shifts relative to oral release of the stop consonant, and on the difference between this relationship and the corresponding one with voiced stops. Roughly:

```
/ti/
mouth   xxxxxxxxxx---------------
larynx  -------------vvvvvvvvvvvv

/di/
mouth   xxxxxxxxxx---------------
larynx  ----vvvvvvvvvvvvvvvvvvvvv
```

Here, xxxxx is oral closure, and vvvv is laryngeal vibration for voicing. VOT is earlier for voiced stops, later for voiceless.

Bruce bn@bbn.com

Date: Tue Jun 23, 1992 3:02 pm PST

Subject: and . . .

I should add that ---- means the mouth or larynx is open, and that considerable VOT differences are perceived only as differences of care in pronunciation (more aspirated t with later VOT, more fully voiced b with earlier VOT), differences of emphasis or stress, or toward the extreme differences in dialect or accent. VOT of both voiced and voiceless stops is earlier for e.g. Spanish /t/ and /d/.

Bruce bn@bbn.com

Date: Tue Jun 23, 1992 5:41 pm PST
Subject: Re: VOT experiment

[Martin Taylor 920623 21:00]
(Rick Marken 920621 and Bruce Nevin 920423 14:07:32)

I've now looked at the Cooper and Nager (1975) article on adaptation of the perception of Voice Onset Time (VOT). I don't think it is very relevant to CSG, though another experiment of a similar kind might be. They had people listen passively (with tongue clenched between the teeth) to 70 repetitions of /i/ or of /r@pi/ (@ is a neutral unstressed vowel called a schwa). Immediately afterward, the subject spoke /r@ti/ in one experiment, /r@pi/ in another. They found that if the subject had listened to /r@pi/ rather than /i/, the voice onset time was shortened, particularly for /t/.

One problem with the study is that the "adapting" VOT for the (synthetic) /r@pi/ was 50 msec, and the natural (after /i/) VOT for the subjects' /p/ averaged 57 msec, 7 of the 20 subjects being shorter than 50 msec. So there wasn't much perceptual adaptation to be expected. And I checked the individual subject data, which showed no relation between their individual VOT average and the direction or amount of shift (If their own VOT was being used as a reference standard that was being affected by the adaptation, one would expect those with long VOT to reduce, and those with short VOT to lengthen after adaptation. That didn't happen).

The only condition that resulted in an appreciable shift of VOT was when they were adapted to /p/ and spoke /t/, which had a "natural" VOT of 75 msec. Adapting to /b/ (VOT 0) and speaking /t/ showed no effect.

Make of this what you will. I don't see any obvious follow-up in PCT terms. I do have problems with the whole "adaptation=fatigue" hypothesis on which the experiment was based. If the synthetic /p/ had had an unnaturally short VOT, but was still perceptually a /p/ rather than a /b/, I would have found the study more interesting, and perhaps there might have been an effect on spoken /p/. But as it is, one has an effect of adapting one phoneme on the production of another, but not on the production of the adapted one. And there seems no a priori reason why the adaptation should have done anything, since the adapting stimulus had a VOT very near the natural VOT for that phoneme.

But one could do an experiment of this kind, usefully, I think. I do not propose to do it.

Martin

Date: Tue Jun 23, 1992 5:53 pm PST
Subject: Re: different cat

[Martin Taylor 920623 21:15]
(Bruce Nevin 920423 08:32:46)

>>As I see it, the key problem is more like: third, how does the brain identify
>>that there exists N objects of a certain kind rather than one strong exemplar
>>of the kind.
>
>One possibly fruitful way to approach this is by attributing
>differentiating perceptions to the individuals in memory and
>imagination.
>...
>Suppose we have ECSs (elementary control systems) controlling for ways
>to differentiate between like perceptions by identifying other
>perceptions associated with them as attributes. Is there any reason
>this is implausible? In such a way we might develop control of category
>perceptions.
>
>One way to track four dots that soon becomes intractible is using their
>history ("the one that started in the northwest corner"). An ECS alert
>for differentia might entertain various hypotheses ("the bounciest one")
>and develop a history of them (". . . is tired now"). A story, not in
>words but in remembered perceptions. ...

I find this approach highly plausible and natural within the HPCT structure. Does it answer your own question? I think it answers mine, which was supposed to mirror yours.

>As we have discussed in the past, something that is structurally
>superordinate is not necessarily superordinate in the perceptual
>hierarchy. Two configuration perceptions going on concurrently, in
>parallel, themselves each provide input to a third configuration ECS.
>The interaction of the two sorts of hierarchy gives the appearance of
>loops on the same perceptual level (we have configurations of
>configurations, idioms are sequences of words, which are sequences).
>There may be loops involving more than one level of the control
>hierarchy. I don't want to open that can of worms now (we're so easily
>distracted), just to put out the hunch that this proposal might provide
>means for dealing with these coherently.

By all means, keep that can of worms closed until you want to go fishing. As soon as you allow same-level loops, you are into the realm in which oscillations and possible chaos lurk. When you loop back to other levels, the demons are tempted to show themselves. I like the terrain, but let's explore the simpler surface first, where loops and like snares are forbidden. I'm sure we will need them soon enough, and, indeed, we have discussed them before.

Beware the Bandersnatch. Martin

Date: Tue Jun 23, 1992 5:59 pm PST
Subject: $$$

[from Joel Judd]

I have to share this nugget of information with people I know will appreciate it:

In the employment section of a large city daily comes this heartening news: for every $10,000 of annual income you expect to get from your new job, you should expect to be unemployed about 1 month. No problem! I'll just apply for unemployment for about--oh--8 1/2 years, and then I'll land the job that will make me a millionare. I knew there was a secret to this job search stuff.

Date: Tue Jun 23, 1992 9:40 pm PST
Subject: VOT experiment

[From Rick Marken (920623.1830)]

Well, I got the JASA article --

Nothing to get excited about; the data is the typical psychology junk -- mostly noise (but what would you expect from psycho-acousticians (hee hee)).

First, thanks to Bruce Nevin for the info about phonemes. Bruce asks:

>Is adaptation a disturbance? (Serious question. I don't understand
>what is going on in situations to which people apply this term.)

I call it a disturbance because if affects the perception of the controlled variable in such a way that action is required to maintain that variable at a reference level. But it could also be seen as something that affects the form of the feedback function ,g(o), that determines how outputs are related to perceptual inputs.

Now, to the Cooper/Nager study.

The basic finding is that adapting to a two syllable word /repi/ leads to an AVERAGE 6 msec DECREASE in VOT. This is not a dramatic effect. Ten of the 22 subjects show NO difference in average VOT or an average INCREASE in VOT. The standard deviation of VOT measures for each subject is bigger (often by a factor of 2) than the average difference in each subject's pre and post adaptation VOT scores.

What this means is that we are dealing with junk data. There is certainly no evidence of a perceptual variable under control. To go off and start modelling based on this data (as the authors do) is simply absurd.

This experiment is a perfect example of the typical psychology experiment; and it shows why, without understanding control, much of the existing data in psychology is almost totally useless to PCT, except by the most remote chance. In this experiment the authors look at the effect of a stimulus variable (the adapting stimulus) on an output variable (VOT). My guess, based on the noisiness of the data (and the highly variable way that the subjects' output relates to the stimulus) that VOT is only indirectly related to whatever a subject controls when speaking the words in the experiment. I have no idea what might actually be controlled in this experiment. The fact that there is an average effect of adaptation on VOT suggests that there is some weak relationship between VOT and a controlled variable. But this is definitely not the way to go about figuring out what the controlled variable might be.

The problem is not just the use of adaptation as a "disturbance". It's that there is no hypothesis about a controlled variable or measures of the quantitative status of that hypothesized variable. I can't imagine starting to build a model of speech until I was able to predict precisely what a person's response would be to every disturbance to the hypothesized controlled variable.

I just got Martin's post on the VOT study.
(Martin Taylor 920623 21:00)

>I've now looked at the Cooper and Nager (1975) article on adaptation of the
>perception of Voice Onset Time (VOT). I don't think it is very relevant to
>CSG,

I think it does have some relevance. It is a great example of why most of the existing psychological data is useless for studying control. This stuff was published in a very prestigious journal. This is the best that psychology has to offer. Believe me, we pretty much have to start over.

>One problem with the study is that the "adapting" VOT for the (synthetic)
>/r@pi/ was 50 msec, and the natural (after /i/) VOT for the subjects' /p/
>averaged 57 msec, 7 of the 20 subjects being shorter than 50 msec. So there
>wasn't much perceptual adaptation to be expected. And I checked the
>individual subject data, which showed no relation between their individual
>VOT average and the direction or amount of shift (If their own VOT was being
>used as a reference standard that was being affected by the adaptation, one
>would expect those with long VOT to reduce, and those with short VOT to
>lengthen after adaptation. That didn't happen).

But they said the effect was even bigger in this experiment than in an earlier study using just syllables for adaptors. I also checked the individual data -- there was a slight positive relationship (r = .34) between subkects' own VOT average and the change -- just the opposite result one would expect if the subjects were controlling perceived VOT by adjusting VOT.

>But one could do an experiment of this kind, usefully, I think.

Do you mean an adaptation experiment or a stimulus/response experiment like this one. I think what they should have done was measured potentially controlled variables -- like spectrograms of the words spoken after adaptation. These could have been compared to spectrograms of words that were not spoken by the subject but picked by each subject the best exemplar of the intended word (after adaptation). At least, that's one possibilty; the idea is to try to see what remains invariant in the input.

I agree that the study of the controlled variables in speech will not be easy -- but the approach taken by Cooper/Nager (which reflects no understanding of the concept of a controlled variable-- and is the approach taken in all psychological experiments in all fields) is not likely to tell you much -- except that stimuli have statistical effects on responses. Well, it does tell you that VOT is probably NOT controlled (though there might be better ways to tease this out). That IS something -- perhaps enough for JASA but not enough for the Journal of Living Control Systems (when it exists).

Regards Rick

Date: Tue Jun 23, 1992 9:40 pm PST

[From Hank Folson (920623)]

Bill Powers [920618] says:
>We can't just go on using the same old customary modes of
>observation, description, and explanation if we want to find the
>significance of the first new concept of human nature since Descartes.

Running through the threads on politics, standards, business, chimps, leadership, etc., I see little consideration of these two basic categories of people:

1. How will a person (living control system) control when s/he does not know that s/he is a living control system, and does not know that everyone else is an independent control system that cannot be controlled without force?

4. How will a person (living control system) control when s/he knows that s/he is a living control system, and knows that everyone else is an independent control system that cannot be controlled without force in a control theory _aware_ world?

Wouldn't the controlling actions be different? In the first case, the actions will be as they have always been in our societies, with all the consequences. There are no Category 4 people yet, because the world is not PCT aware. Here are three transitional categories:

1A. (A small but important sub-category) How will a person (living control system) control when s/he does not know that s/he is a living control system, and does not know that everyone else is an independent control system that cannot be controlled without force, but has through their life experience developed a Systems Concept and controlling techniques that are compatible with PCT?

2. How will a person who has lived most of their life and had most of their experience in life in a Category 1 world as a Category 1 person, but then founded/pioneered/studied/was seduced by PCT, attempt to understand and control in our Category 1 world? Will s/he only use the new paradigm, or will s/he drift back and forth between paradigms?

3. How will a person who has been raised in a Category 1 world, and whose Systems Concept is rooted in Category 1, and who may even have a vested interest in the world remaining a Category 1 world, control (still in determined ignorance of their being a control system) when faced with the PCT concept of the world?

Why have I defined these Categories in terms of questions? In part in response to Bill's statement quoted above. In part because we Category 2 people must remain aware of our Category 1 roots, and that the world will stay Category 1 for an unpredictable time, and most of our controlling will be with Category 1 and 3 people and their worlds. This need for awareness is also why I chose not to use the word 'behavior' in the definitions, only the word 'controlling', so we will be more likely to concentrate on what others are controlling for, and not waste time discussing behaviors (in the traditional S-R sense of the word). If we don't keep the questions in our minds, won't we drift into ineffective modes, make honest errors in analyzing and controlling and so on? After all, we are all simply living control systems, and there is nothing in control theory that says just because you know control theory you will control in the most effective way, because everything is all perception. If you lose sight of what category you are in and what category others are in, will you control effectively for what you want from others?

Would you gain any new insights by using these categories were you reread and analyse again what is going on in the recent posts on the newspaper organization, leadership, Chimps, the Jury comments?

Categorically, 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 Jun 24, 1992 5:50 am PST
Subject: what's the difference?

[From: Bruce Nevin (Wed 920424 08:22:46)]

I suspect that there is more than one sort of thing being called adaptation. I have proposed before that one of these is a change in the value of a reference signal, basing it on short-term memory more than on long-term memory. This would be real adaptation. Another sort of thing might be reduction in gain. Later repetitions of a word in discourse are produced with less careful articulation than the first occurrance, Martin tells us (a fact that, to my delight, fits very well with the Harrisian notion of reductions). With a reduction in gain there might be a statistical tendency but not the individual consistency that is the hallmark of control.

There is always the possibility of boredom, listening to 70 repetitions of a synthesized while (whilst) clenching ones tongue between one's teeth.

The idea of adapting the reference signal has an interesting twist for language learning. We do accomodate small shifts of dialect without noticing, preserving contrast of words; we also resist large dialectal differences at a different level of control, the contrast between the kind of person who speaks noticeably that way and the kind of person with whom I am most familiar. If it were possible to vary the relevant parameters for a difference in pronunciation (say English vs. Spanish VOT) in ongoing dialog with the student, the student might "adapt" to the new reference values in graduated steps, never encountering a differential sufficient for control on the self-image level. One can imagine a SF world, where the student was immersed in an environment with androids whose reference signals were able to be tuned gradually in this way. I'm not sure if it would be possible for a human instructor to make the gradual adjustments in a consistent way--sensitive always to the student's current and changing settings of those signals, and while carrying on natural and engaging conversation.

(Martin Taylor 920623 21:15) --

Replying to me (920423 08:32:46)

>>One possibly fruitful way to approach this is by attributing
>>differentiating perceptions to the individuals in memory and
>>imagination.
>>...
>>
>>Suppose we have ECSs (elementary control systems) controlling for ways
>>to differentiate between like perceptions by identifying other
>>perceptions associated with them as attributes. Is there any reason
>>this is implausible? In such a way we might develop control of category
>>perceptions.
>>
>>One way to track four dots that soon becomes intractible is using their
>>history ("the one that started in the northwest corner"). An ECS alert
>>for differentia might entertain various hypotheses ("the bounciest one")
>>and develop a history of them (". . . is tired now"). A story, not in
>>words but in remembered perceptions. ...

>I find this approach highly plausible and natural within the HPCT
>structure. Does it answer your own question? I think it answers mine,
>which was supposed to mirror yours.

Yes, I hope so, and thank you for taking up the question and rearticualting it so well.

Now, what might a difference-detector ECS accept as input? Two signals from an ECS, one associated with one set of signals from other ECSs, another associated with an intersecting but different set. A discontinuity in the signal or signals that stay the same, reflecting a discontinuity in sensory input or a shift of attention. (Attention involves at least the real or imagined focus of sensory organs on a target and the adjustment of gain on selected ECSs, as for example tuning out extraneous noise for conversation in a crowd.) The discontinuity marks a shift from one individual to another, unless the difference-detector itself reports continuity (no differences that make any difference).

```
          a  m  p                                a  b  c
           \ | /                                  \ | /
        b -- X -- q                            m -- X -- n
           / | \                                  / | \
          c  n  r                                p  q  r

  Individual    Individual              Individual    Individual
  A = X plus    B = X plus              A = X plus    B = X plus
  {abcmn}       {pqrmn}                 {mabcpqr}     {nabcpqr}
```

Here are two hypothetical cases of "detecting an X" plus associated perceptions {abcmnpqr} that may be present. X is a category perception.

My proposal regarding category perceptions is that they develop from the association of e.g. {abcmn}, {pqrmn}, {mabcpqr}, {nabcpqr} as distinct individuals by a difference-detector. The difference-detector then constitutes the category perception out of these perceptions of attributes. Any association-set like {abcmn} or {pqrmn} (or in the second hypothetical case, any association-set like {mabcpqr} or {nabcpqr}) indicates that an X is present. Were one to define an X, one would say anything that has attributes {mn} (in the first case), with other attributes {abc} or {pqr} optional. (In the second case, the defining attributes would be {abcpqr} and the variability would be between m and n.)

According to this proposal, there are difference-detectors, and there are associations of perceptual signals (attributes) in associative memory, but there are no category detectors per se. Does this still sound feasible and sensible?

Bruce bn@bbn.com

Date: Wed Jun 24, 1992 8:00 am PST
Subject: Meanwhile, on Another Network...

Dear Readers,

I would be very interested in hearing from anyone who has attempted to create a model of an antoganostic nuromuscular system using matlab or simulab. I have been trying to recreate the work of Stark and his co-workers without much success. I most interested in the elbow joint, but any joint would be a start.

Any input or references in this direction would be greatly appriciated.

Thanks,
Jim Fee
Research Engineer
A. I. duPont Institute.

The following note may be of interest to the readers of the forum:

The McDonnell-Pew Program in Cognitive Neuroscience is accepting proposals for support of research and training in cognitive neuroscience. Preference is given for projects that are not currently funded and are interdisciplinary, involving at least two areas among clinical and basic neurosciences, computer science, psychology, linguistics and philosophy. Research support is limited to $30,000 a year for two years. Postdoctoral grants are limted to three years. Graduate student support is not available.

Applications should be postmarked by August 1, 1992 to:

Dr. George Miller
McDonnell-Pew Program in Cognitive Neuroscience

Date: Wed Jun 24, 1992 8:18 am PST
Subject: Re: VOT experiment

[Martin Taylor 920624 10:45]
(Rick Marken 920623.1830)

On the VOT study
>I agree that the study of the controlled variables in speech will
>not be easy -- but the approach taken by Cooper/Nager (which
>reflects no understanding of the concept of a controlled variable--
>and is the approach taken in all psychological experiments in
>all fields) is not likely to tell you much -- except that stimuli
>have statistical effects on responses. Well, it does tell you that
>VOT is probably NOT controlled (though there might be better ways
>to tease this out).

Leaving aside Rick's obligatory rant--how could one have expected Cooper and Nagel to have understood the concept of a controlled variable?--I don't understand how Rick could possibly come to the conclusion that VOT is probably NOT controlled. That conclusion would lead one to believe that people are unable reliably to distinguish /b/ from /p/ in production (at least the allophone without the burst). VOT is clearly perceived, even, apparently, by chinchillas and newborn humans. It is a major contrast in the perception of language, and I seem to remember even that some languages contrast it in three levels rather than two (don't quote me on that). Certainly some languages place the VOT boundaries at different timings from others, so the discrete VOT timing contrasts are not quantitatively innate. I find it highly unlikely that VOT is not a controlled percept.

The Cooper/Nager study is very strange, whether you are of the PCT persuasion or not. I do not understand why any adaptation effects relating to VOT should have been observed, unless they had to do with variability, since the adapting VOT was very close to the natural VOT for that phoneme. Nevertheless, they did observe a change in the VOT for the production of another phoneme, 18 of the 22 subjects showing a decrease. (I'm not going to be sidetracked into another "stale and unprofitable" discussion of statistics at this point--but another day, perhaps, we can heat it up in the microwave).

As for things being printed in prestigious journals--there's lots of garbage printed in all journals. Whether something good is rejected or something bad is printed depends so much on one or two individual reviewers and whether they had a good breakfast the day the looked at the paper. I always try to have a first read to get an overall impression, and then set it aside for a couple of weeks to see whether a detailed look confirms that impression. But I think many reviewers just skim it and say "OK" or "junk" without checking
the data or the analyses. I would not have recommended publishing Cooper/Nager
because of the mystery about the 50 msec VOT of the synthetic /p/, as well as for other reasons.

Rick's suggestion about studying the spectrum seems less profitable than studying the VOT. I am under the impression that the spectrum associated with a phoneme is not a controlled variable, since the perception of it is highly context-dependent. One can measure the physical acoustic spectrum, but that doesn't necessarily tell you about the perceived timbre.

It's all very well to rail on about measuring stimuli and responses, but they are at least measurable, even if the measurements don't tell you much that's useful. To measure what is controlled, you have to intuit (empathize, introspect, anthropomorphize) some complex that can be computed from physical measurements, disturb it or the perception of it, and see whether the disturbance is resisted by further measurements on the "same" variable (and "same" is an issue in itself). If you intuit wrong, you will get partial or no control, and more or less statistical variability. That wrongness may be in your intuition as to what is controlled, or in the effect of your disturbance on what is controlled. My hunch is that in the Cooper/Nagel case the wrongness is in the perceptual disturbance more than in the variable.

Martin

Date: Wed Jun 24, 1992 8:23 am PST
Subject: what's the difference?

[From Rick Marken (920624.0900)]

Bruce Nevin (Wed 920424 08:22:46) says:
                    ^ (where the hell is BBN. It's JULY on the west coast)

>I suspect that there is more than one sort of thing being called
>adaptation.

That's for sure. I think what the VOT people were trying to do was adapt in the sense of fatiguing a sensory system. Adaptation is a well known perceptual effect. Probably the best known are color aftereffects. Remember the green and orange striped american flag in your introductory psych text. When you stare at it for a while (1 min say) you are presumably fatiguing one side of an opponent process color system -- the green is fatiguing the green side of a red-green system. If you look at white paper after this adaptation (white has both "red" and "green" in it) the green system is all tuckered out -- so it can't fire as strongly as the red side, and it can't inhibit the red side as strongly as the red can inhibit the green. Hence, you see ol' glory in red white and blue after adapting to the green, orange and yellow (I think? -- I can't remember color compliments).

This kind of adaptation also happens with motion (the "waterfall" effect). I think Bill Powers was actually helping with a control study based on this kind of adaptation. If you stare at a waterfall (downward retinal motion) for some time, then look at the stationary cliff next to the waterfall, the cliff seems to move UP.

These adaptation effects are a nice substitute for hallucinogenic drugs -- and they also are also a nice tool for studying perceptual processing. I suppose that the VOT people imagined that there was detector for VOT. But I don't know how they thought the adaptation might work. I guess there would have to be some idea what the "neutral" value of VOT is (like the stationary cliff in the waterfall effect and white in the color effect). I suppose zero VOT would be the neutral value? Anyway, after adaptation to a VOT of +X, a zero VOT should then sound like -Y. It's not clear what the effect should be if you present a VOT near the adapted value (as they did in the study). In the color and waterfall effects, there would be virtually NO effect. That is, if you adapt to green and then show green, it still looks green. If you adapt to "up movement" and then show up movement it still looks up.

Maybe they just did the VOT study wrong. If there really is a VOT detector that adapts, then if you adapt the detector to a +58msec VOT then a sound with zero VOT should sound like it has an approximately -58 VOT. So to produce the sound correctly, the subject would have to delay the VOT (which is usually zero) by about 58 msec.

Such a study should be done quantitatively (varying the degree of adapting VOT over some range and watching for quantitative opposition by the speaker. If the opposition is not perfect (correlation between adapting VOT and spoken VOT = .99 or greater) then keep looking for the controlled variable.

Best regards Rick

Date: Wed Jun 24, 1992 10:10 am PST
Subject: Chimps; Xs; HPCT as tool

[From Bill Powers (920623.1100)][Delayed]

Martin Taylor (920622.1930)]

You and Bruce are right that the Test requires anthropomorphizing. But Sagan and Drayer did it wrong in two ways.

First, they were crossing species. Instead of using the lowest common denominator of description, they used the highest: they described the (potentially controlled) perceptions in words appropriate for human beings. I would pick terms that imply the least sophistication of perception and control, not the most, in making my guesses about another species.

Second, the description of the social interactions employs stereotypes that probably don't reveal what even human participants are controlling for in such situations. There's a bit of cynicism in the description: I'll bet Carl Sagan never stood in a crowd and reached out to touch a charismatic leader like George Bush, or that he ever bowed down in awe before any other person. That's just his (their) stereotype of how the common masses behave. If S&D could get inside the heads of most "common" participants in rituals and ceremonies, they might be shocked at finding how much of the same sort of cynicism is there. If Sagan could hear people reacting to him when he strikes an imposing pose on television, he might be furious, as well.

Most people, I think, behave in ritual situations as they do not because they have some corresponding inner conviction, but because they think everyone else does. You take off your hat, stand, and put your hand on your heart when the anthem is played and the flag is raised at Yankee Stadium, because a show of patriotism is what's expected -- even if you're Canadian or an anarchist. The main pressure is concerned with controlling for appearances, not beliefs. If you act acceptably, you can believe anything you want. Your real goals are your own business.

A great many of the silly and boring things that people do become understandable if you know that lots of people are controlling for stereotyped actions that they see others carrying out, simply to conform with what's expected in the situation. I don't think there are nearly as many True Believers as there seem to be.

When we find out what people are really controlling for, I don't think we will find much relationship to commonly-held stereotypes like those that Sagan and Drayer described (or imagined). The most commonality will be found in the overt behaviors that other people can see, and by which individuals are judged in social circumstances. This could even be true in a chimp society. The gestures, if produced simply because they're expected, don't have much to do with anything except the penalties for nonconformity. Inner and outer conformity, however, are not the same thing.

>I don't find CSG revolting, and I don't want to. Jolting is OK, but
>revolution seldom has good results, politically or scientifically.

Seldom isn't never. There have been some pretty good revolutions among the bad ones.

>There's lots of good food for thought out there among the garbage. Some just
>needs to be made a bit more tasty by being taken with a grain of salt.

Salted garbage is still garbage. I don't want to sort through garbage for something that looks edible, even if it's there. When things get messy enough, it's time to go to the grocery store and start over, even though, in principle, you could sort out the peanuts from the coffee grounds.

RE: perception of multiple Xs

>Bill points out that in principle one can only [track]the objects if
>the targets maintain differences from the distractors in at least one
>of position, velocity and acceleration. That's quite true, and
>experimentally verified. But even after two have been confused, the
>observer can still track about four (depending, I think, as Bill says,
>on the motion statistics of the ensemble). One of the four may be,
>from the experimenter's viewpoint, wrong, but the observer sees the one that
>best fits whatever tracking method is used by the ECSs that are employed.

I don't understand the "But...". In the "gating" model I proposed, the tracking systems don't give up when they confuse two objects. They go right on tracking the n objects, as if the original ones were still being tracked. The problem is that the systems are NOT confused. If four objects were being tracked, four continue to be tracked. They are not necessarily the original set, but the tracking systems don't know that. If a given tracker knew that there was more than one item in the gated area, it wouldn't confuse two items. Such gated tracking systems simply center the centroid of whatever is in the gated area, on the ASSUMPTION that it's a single object.

>As I see it, the key problem is more like: third, how does the brain
>identify that there exists N objects of a certain kind rather than one
>strong exemplar of the kind. I've forgotten how Bruce originally put >it,
but perhaps I can paraphrase. If there is an ECS controlling for >the
perception of an X, it should satisfy its reference if an X is in >its
input. What distinguishes thes existence of exactly three Xs in >its input
from the existence of one X.

If the ECS is controlling for the perception of "an X", then any number of Xs greater than zero at its input will result in zero error. If it is controlling for "exactly n Xs", then there are two ECSs involved: one elementary control system can handle only one variable. The second ECS is controlling for number of objects. Another level is required to handle the "exactly" part: the number must be neither less than nor greater than n, AND for each item taken alone, the item must yield the perception of X rather than Y. This is a logical condition -- at least a relationship (the AND relationship).

In the pandemonium model, the presence of an X anywhere in the visual field, with any size or orientation, is enough to produce perception of X-ness. I proposed several posts ago a model in which each instance contributed to the total X-ness. That was wrong -- that would mean that X-ness is not independent of the attribute of position. All that can change the amount of perceived X-ness is that some or all items, while still recognizeably X, differ from the canonical X to some degree. The number of items in the visual field brings in a new dimension of perception.

If you're presented with a visual field containing dozens of Xs and Ys, and are asked to find 3 Xs, you will be able to do so, but only by serially searching among the items and discarding those that aren't Xs. If there is no contiguous group of 3 Xs, you probably won't be able to construct a perception of 3 Xs, easily or perhaps at all. You'll have to handle it by counting at the symbolic representation levels, putting gates around qualifying objects, and so on.

If all the items are Xs, then the X-ness perception is satisfied wherever you look, and the 3-ness perception is satisfied by any configuration of 3 objects. If explicit perception of position is introduced (it doesn't have to be), we could also require that "there-ness" be perceived - i.e., a specific location or set of locations, either predesignated or specified by relationship ("not the same").

For any of these ideas to work, I think we have to assume that judgements are made within some limited region around the center of vision, or attention, or both.

While this still doesn't handle all problems, it does help us to remember that an ECS perceives ONLY ONE PERCEPTION. In describing any experimental situation in words, it's easy to pack multiple dimensions of perception into unitary-seeming sentences or phrases, like "three Xs." n-ness and X-ness are independent dimensions of perception, and can vary independently. To see "n objects of a given kind" is to see (1) n objects, regardless of kind, and (2) a given kind of object. If the additional condition is set that (3) no OTHER kind of object is to be perceived, another dimension of perception is introduced. If the condition is that (4) n and only n objects shall be perceived, we have still another condition to satisfy and perceive, which would be different from the condition (4a) at least n objects, or (4b) at most n objects. Also, "see n objects and see X-ness" is a reference condition different from "see n perceptions such that each is an object AND is an X."

When you have to analyze a perceptual situation down to individual and independent perceptions, a lot of ambiguities are usually revealed. This is one of my beefs with psychological experiments: most of them are ambiguous in ways their authors overlooked because of taking informal verbal descriptions for granted. For the same reason, most such experiments involve phenomena that are immensely more complex than the authors recognize -- this is one reason for the lousy correlations that are published.

>Bill said in some posting shortly after I joined this group that >whenever he had thought of a level-jumping control, it turned out not >to be (appropriate/correct/necessary/simple?). Maybe so. In the >neural network business, multilayer perceptrons often work well, but >there are other architectures that are more appropriate for complex >problems, including specifically gated modular architectures in which >smallish modules solve sub-problems, and gating structures determine >which subnetworks present their solutions to higher modules.

See Gene Boggess' excellent post from 920622. If all you want to do is solve a problem, there are many ways to do it. But if you want the solution to be consistent with neural architecture and function, behavioral facts, and subjective experience, the possible choices are greatly reduced.

In BCP and elsewhere I said that level-jumping perceptions seem to be permissible, but jumping levels in the downward direction leads to a problem. The reason in general is that a control system receives its net reference signal from multiple sources. If a system of level n+2 contributes a reference signal to a system of level n, the result will be a disturbance of all systems of level n+1 that are using the same level-n system to achieve their own goals. Those systems will react by cancelling the effect of the level n+2 output. Furthermore, the level n+2 system does not perceive in the right terms for attaining a compromise with the competing level n+1 systems: it would be like looking for a simultaneous solution of a set of equations in which the variables are x,y,z, and "consistency".

This problem doesn't exist in every possible case; it is obviated when intervening levels don't come into play -- when, for example, there are transitions under control, but no specific events. Then you can control relationships among transitions.

>Perhaps control nets might work better on complex problems if conflicts
>can be resolved by gating structures that permits some modules but not
>others to exercise control?

This is already part of the HPCT model, although not all MODES of control are "official" yet. Particularly at level 9, logic and programs, it's possible for control systems to work by sending reference signals to some lower systems but not others: select one sequence rather than another, or where sequence is not controlled, one category of action rather than another. And of course even within the program level it's possible to have nested subroutines -- why not? That's just a matter of what program is running, not of basic organization. Anything that can be programmed in a neural computer can run at the program level. There's nothing about any particular program that characterises human nature. I wouldn't be surprised if ALL propositions at this level can be found exemplified in SOMEONE'S head -- if only the head of the guy who thought up the proposition.

>Simplicity is to be preserved where possible in science, and HPCT is a
>simple structure. If it can solve problems as basic to perception as
>"there's an X and there's another" or "I see a lot of Xs among the Ys"
>without the introduction of new structures or concepts, so much the
>better. At present I don't see the answer.

HPCT is expressed in terms too general to provide specific answers to specific questions -- particularly when they are questions about the content, rather than the form, of behavior. HPCT is really a toolkit. The claim is that by using the elements of this model, by making specific propositions that employ only these elements, one can construct explanations of all behaviors and organizations of behavior. The only way to test this claim is to see whether the design principles and parts list in the HPCT model will suffice, given enough creative effort. Of course if the effort involved is greater than the effort needed when using some other kit of ideas, we would become suspicious of the HPCT model and either modify it or abandon it.

This is why I can't answer a lot of questions, like "What does HPCT have to say about cognitive dissonance?" It doesn't have anything to say until you go through the effort of creating something to say using the concepts of HPCT instead of whatever concepts underly "cognitive" and "dissonance" (if any). All the problems of behavior that come out of other approaches are defined in terms of a different conceptual toolkit and a different assumed parts list, usually irrelevant to the concepts and components of HPCT.

Tho most difficult questions to answer are those that contain implicit theoretical assertions, or unconscious assumptions, or taken-for-granted interpretations. I've been asked, for example, "What causes anxiety?" Such a question really stumps me at first, because not only can't I answer it as stated, but I doubt whether it HAS an answer as stated.

Under HPCT we don't deal in things that "cause" things in a simple 1-2 manner. Terms like "anxiety" don't refer to any objective state of a system, but to how a person perceives an internal state. An internal state of what? Evaluated with respect to what reference level? And of course in asking for a "what" that causes anxiety, one is asking for a force that, acting on a system, makes it respond in a specific way -- and in the HPCT model, that never happens. The real answer to a question like "What causes anxiety" is "Nothing. There's no thing called anxiety, and there's no cause for things that don't exist, and anyway 'causation' is the wrong concept." And then, of course, to handle the PHENOMENON with which this person is concerned, one would have to start from scratch and study it as a phenomenon of control. It's a common predicament that for me to answer certain classes of questions, I'd have to do 10 years of research in HPCT terms (which the asker wants me to report on instead of doing it himself or herself).

A basic premise of HPCT, that an elementary control system handles only one variable in one dimension of variation, forces us to analyze phenomena into single variables of just a few classes. Doing this usually reveals that we have posed more than one question, thinking it was a single question. It doesn't really matter whether the premise here is absolutely true all of the time; just thinking that it's true is often enough to lead to reformulation of a problem in a way that suggests multiple solutions.

Gene Boggess (920622) --

>Unfortunately, not much AI is actually based on ANY model of human
>intelligence; "whatever works" seems to be the watchword. So I don't
>know that HPCT will have an overwhelming impact on the field; but it
>could, and it should. It is certainly affecting how I think about my
>own AI work.

A lovely post -- so THAT's what's been going on! I think a lot of people would be interested in your description of just what changes in your approach to AI took place. I expect that many people on this net have suffered through a paradigm shift like yours, and will have a deep understanding of the struggle and the cost involved.

Best to all,   Bill P.

Date: Wed Jun 24, 1992 10:19 am PST
Subject: Psychological research

[From Rick Marken (920624.1030)]

(Martin Taylor 920624 10:45)

>I don't understand how Rick could possibly come to the conclusion that VOT is
>probably NOT controlled.

I guess it was the total absence of any evidence that VOT is controlled.

> That conclusion would lead one to believe that people are
>unable reliably to distinguish /b/ from /p/ in production (at least the
>allophone without the burst). VOT is clearly perceived, even, apparently,
>by chinchillas and newborn humans. It is a major contrast in the perception
>of language

VOT is something that a bunch of phoneticians think is important (like the
biologists who think the fall of the bat is important). It is just a fairly
noticable part of the perceptions of the observer (of speech or bat), but is
not necessarily a variable that is being controlled by a speaker (though
what is controlled may be partially a consequence of VOT -- just as what
the bat is actually controlling is a partial consequence of its downward
acceleration).

>Nevertheless, they did observe a change in the VOT for the production of
>another phoneme, 18 of the 22 subjects showing a decrease. (I'm not going
>to be sidetracked into another "stale and unprofitable" discussion of
>statistics at this point--but another day, perhaps, we can heat it up in the
>microwave).

I think it's time for this sidetrack because my main objection to the JASA
VOT study is the extraordinarily poor quality of the data. Yes, there were
average effects. But look at the VARIANCE. Each subject adapted and pronounced
the target word 20 times without adaptation and 20 times with adaptation.
The standard deviation of the VOT measures over these two sets of 20 trials
for
each subject was on the order of 10 msec. That's the AVERAGE deviation of
what is presumed to be the controlled variable (VOT). Thus, the presumed
controlled variable varies over repetitions with NO DISTURBANCE more than
its average increase with the disturbance (adaptation) -- 10 msec vs 6 msec--
not my idea of good control over VOT.

I am not faulting the people who did this study for not knowing PCT or
for doing poor research. I am saying that this is the best research that
can be expected from the S-R statistical approach. This kind of dreck is
tolerated because nobody expects to get better results than this anyway.
And indeed they can't get better results if they look at behavior from
an sr perspective (the perspective demanded by this kind of research) --
people have tried to get better results for decades and this is the best that
they can do. The reason, of course, for the poor results is that people
are not sr devices so the best you ever get is statistically significant
relationships between s and r.

Martin seems to think that the problem with this study is that they didn't
do the adaptation right. I think that that is irrelevant; this study
is useless because it is all noise. As I said in an earlier post, if

the relationships in your data are not .99 or greater then you should try to fix the research until you get such relationships. This VOT research should have gone into the wastebasket -- and the reserachers should have kept trying until the subjects made exactly the response expected on every trial in every condition.

>It's all very well to rail on about measuring stimuli and responses, but
>they are at least measurable, even if the measurements don't tell you much
>that's useful.

I am "railing" about trying to make sense out of random noise. This VOT study is a great example of a study that purports to have discovered something (the importance of VOT in speech perception and production) when it has done nothing of the kind. It is just flat out misleading -- just like all statistical psychological research. And I consider that a disservice to humanity (because it is what I consider to by LYING). I rail because it makes me mad when people mislead other people and claim the mantel of science to give force to their mendacity. Rail.Rail rail.

> To measure what is controlled, you have to intuit (empathize,
>introspect, anthropomorphize) some complex that can be computed from physical
>measurements, disturb it or the perception of it, and see whether the
>disturbance is resisted by further measurements on the "same" variable (and
>"same" is an issue in itself). If you intuit wrong, you will get partial or
>no control, and more or less statistical variability. That wrongness may
>be in your intuition as to what is controlled, or in the effect of your
>disturbance on what is controlled. My hunch is that in the Cooper/Nagel case
>the wrongness is in the perceptual disturbance more than in the variable.

The test for the controlled variable starts with intuition and guesses but it is a quantitative process (you MEASURE the controlled variable) and you don't stop the process (or report success) until you find disturbance/action relationships at the .99+ level. It's a whole new approach to research; with all new standards of excellence -- because you know what is going on (control) and you know how to detect it (the test). So if your results are not nearly perfect, you know you have done something wrong (or not correctly identified the controlled variable). You don't need to report junky results (unless you have to publish or perish) because they are useless (except to let others know of blind alleys -- so I guess the literature of psychology could be considered a roadmap to one giant blind alley). It's not a matter of getting less statistical variablility -- it's a matter of getting NONE. We know this can be done because we know that control is not a statistical process (how much variability do you tolerate when you control the you balance) and we have done it (see my book and the studies by Tom Bourbon and Bill Powers)

In the Cooper/Nagel study the wrongness was in the criterion for what is an acceptable level of precision for a scientific result. Psychological research ain't going to get nowhere as long as it continues to accept results like those of Cooper/Nager as being anything other than useless noise. It's time to grow up, psychology.

One last point. Some might say -- "well statistical research is the best we can do now. It's better to have some low quality data then NO data ay all". I say bull. What we know from low quality data is, in my mind, useless or, worse, misleading. And in psychology this can actually hurt people -- if they believe the statistical results that are published. I think it will take some time before PCT is able to start developing a body of high quality data that can start answering our psychological questions authoritatively. Until then, I think the best answer to questions like "what do people control when saying things like /reti/?" is "I DON'T KNOW". This is generally good advice to all people who are trying to make sense of their world -- religiously or scientifically -- and not trying to impress their little brother or sister. Try it. It's very relaxing. Is there a god? I DON'T KNOW. Are my standards the best? I DON'T KNOW. Is VOT a controlled variable? I DON'T KNOW (but probably not).

Regards Rick (Microwave) Marken

Date: Wed Jun 24, 1992 10:31 am PST
Subject: July is bustin' out all over

[From Rick Marken (920624.1032)]

In my earlier post to Bruce Nevin I wrote:

```
>Bruce Nevin (Wed 920424 08:22:46) says:
>                    ^ (where the hell is BBN. It's JULY on the west coast)
```

Martin Taylor just informed me that it is still only June in Toronto. A quick look at my Timex confirms that it is still June in LA as well. My mind was in July because I am beginining a family vacation to London on July 1 and I'm just finishing up stuff at the office in preparation.

So I go with Martin on this one; it is, indeed, still June (darn).

Best regards Rick

Date: Wed Jun 24, 1992 10:38 am PST
Subject: Re: what's the difference?

[Martin Taylor 920624 13:00 (still June in Toronto)]
(Rick Marken 920624.0900)

```
>Bruce Nevin (Wed 920424 08:22:46) says:
>                    ^ (where the hell is BBN. It's JULY on the west coast)
>
>>I suspect that there is more than one sort of thing being called
>>adaptation.
>
>That's for sure. I think what the VOT people were trying to do was
>adapt in the sense of fatiguing a sensory system. Adaptation is a well
>known perceptual effect. Probably the best known are color aftereffects.
>Remember the green and orange striped american flag in your introductory
>psych text. When you stare at it for a while (1 min say) you are presumably
>fatiguing one side of an opponent process color system -- the green is
```

>fatiguing the green side of a red-green system. If you look at white paper
>after this adaptation (white has both "red" and "green" in it) the green
>system is all tuckered out -- so it can't fire as strongly as the red side,
>and it can't inhibit the red side as strongly as the red can inhibit the
>green. Hence, you see ol' glory in red white and blue after adapting to
>the green, orange and yellow (I think? -- I can't remember color
>compliments).
>
>This kind of adaptation also happens with motion (the "waterfall" effect).
>I think Bill Powers was actually helping with a control study based on this
>kind of adaptation. If you stare at a waterfall (downward retinal motion)
>for some time, then look at the stationary cliff next to the waterfall,
>the cliff seems to move UP.
>...
>That is, if you adapt to green and then show green, it still looks green.
>If you adapt to "up movement" and then show up movement it still looks
>up.

Oh, wow...disinformation piled on misinformation! I know that many people think of aftereffects as the consequence of fatigue, but they can't be, at least in most cases and perhaps in all.

Disinformation first: It is NOT true that if you adapt to up movement and then show up it still looks up. The aftereffect ADDS to the real movement, so that you can get an impression of zero movement by presenting the correct amount of up movement. Things obviously appear continuously higher, but there is no apparent movement. The perception of movement is separate from the perception of change of position. I did lots of experimental and theoretical work in this area in the 60's, and one thing that subsequent work cannot change is that the phenomenology is very complex. A naive view based on "adpaptation=fatigue" cannot work because it predicts a lot of phenomena wrongly. "Adaptation=improved perceptual precision" accounts for quite a few of the phenomena that "adaptation=fatigue" does not, and predicts numerically as well as qualitatively. In the red-green case, it is possible (even likely) that fatigue plays a part. In the movement case, it is less probable, and when we come to the more shape-based aftereffects, it is not likely at all.

What control study was Bill doing? Maybe Bill can answer. I did several studies around 1962-3 in which I had people counter the after-effect of motion,
to measure its time course. I asked them to keep the perceived motion at zero after a period of viewing the moving thing. The decay of the after-effect was double exponential, quantitatively proportional to the square root of the duration of the adaptation. There are at least two effects there. The neutralization of the adapting percept, on the other hand, is logarithmic in the square root of time for at least 10 minutes of observation, and perhaps much longer. It is different again. (Reference examples: Perceptual and Motor Skills, 1963, 16, 119-129; 1963, 16, 513-519; 1964, 18, 885-888; Perception and Psychophysics, 1966, 1, 113-119).

I don't follow Rick's analysis of the VOT detector system at all:

>Maybe they just did the VOT study wrong. If there really is a VOT detector
>that adapts, then if you adapt the detector to a +58msec VOT then a

>sound with zero VOT should sound like it has an approximately -58 VOT.
>So to produce the sound correctly, the subject would have to delay the
>VOT (which is usually zero) by about 58 msec.

If there really is a VOT detector, it asserts whatever VOT is appropriate for the phoneme in question as a reference. Zero VOT has no special significance, and may not even be auditorily identifiable outside its effect on the categorization of a phoneme. Even if zero were a special reference point for VOT, as "vertical" is in vision because of gravity, one would expect adaptation to +50 msec to have an effect no greater than about 5 msec on the perception of zero. And that 5 msec would be different for other non-zero values of VOT, being probably (though not necessarily) of different sign in different regions of absolute VOT.

Bruce is right. There are several different effects of adaptation. Adaptation
is a procedure, not a "perceptual effect". The adaptation procedure can lead to fatigue, to excitation, to changes in the relative precision of perception at different places on a continuum, to the diminution of deviations from a reference (independent of fatigue), and to who knows what else. These in turn lead to perceptual effects, which may differ accordinag to circumstances, for any of the mechanisms.

Rick often proclaims that all psychological studies done outside the control paradigm are worthless. This shouldn't give him the right to assert his own view of the world that they study, in contradiction to the results they obtain. You can't claim better truth by throwing away data than by looking at what data you have, however wrongheaded the data gatherers might have been. A solipsist world is not the domain of science.

>Such a study should be done quantitatively (varying the degree of adapting
>VOT over some range and watching for quantitative opposition by the speaker.
>If the opposition is not perfect (correlation between adapting VOT and
>spoken VOT = .99 or greater) then keep looking for the controlled variable.

Yes, one would have to do soemthing like that, if we could rig a vocoder-like system to modify the VOT of the subject's own productions. But since at least part of the sensation leading to the VOT perception is likely to be from internal sources (feeling the vibration in the throat, for example), it would be hard to expect such high correlations even if one could devize a real-time VOT-shifter.

Sorry to be so harsh, Rick, but that message really seemed wrongheaded, if not bull (headed).

Martin

Date: Wed Jun 24, 1992 11:24 am PST
Subject: Re: Chimps; Xs; HPCT as tool

[Martin Taylor 920624 13:50]
(Bill Powers 920623.1100)

I'm sorry my language is so obscure. I was really agreeing with you that the tracking system did not get confused when two objects came close and then departed. It is the experimenter that says that the tracking system was confused because the item tracked after the "collision" was in his/her view different from the one being tracked before.

>If you're presented with a visual field containing dozens of Xs and Ys, and
>are asked to find 3 Xs, you will be able to do so, but only by serially
>searching among the items and discarding those that aren't Xs. If there is
>no contiguous group of 3 Xs, you probably won't be able to construct a
>perception of 3 Xs, easily or perhaps at all. You'll have to handle it by
>counting at the symbolic representation levels, putting gates around
>qualifying objects, and so on.

Ah...a comment that I can disagree with, which might further establish the question (but you can note that I believe the problem has already been solved by Bruce).

There are two ways that the human perceptual system can solve this problem, and it depends on whether the distinction between X and non-X can be categorized
in some "natural" (as yet undefined) way. If the Xs are categorically different from the non-Xs, it is likely that there is no sequential search process, and that the time to develop the percept is independent of the number of non-Xs. But if there is no such "natural" category distinction, the time to develop the percept will increase substantially with each additional non-X (substantial means a few tens of msec per item). It is not true that you have to "gate" each X in turn for counting. I have no idea to what number of Xs this applies. ("Natural" seems to mean some kinds of shape contrast, or at a higher level vowels versus consonants or numerals against letters. Perhaps other categories work as well. There was a suggestion
that function words against content words had the same effect. I don't know where this all ends.)

> If all you want to do is
>solve a problem, there are many ways to do it. But if you want the solution
>to be consistent with neural architecture and function, behavioral facts,
>and subjective experience, the possible choices are greatly reduced.

Yes, I fully agree, and that is what I am trying to do by presenting various behavioural and phenomenal observations and asking how the simplest models deal with them. I'm glad to know that "official" HPCT includes gating effects.
I'm in the process of producing a note (for CSG-L, but long) arguing that the degrees of freedom constraints dictate such gating, possibly at quite low levels, and that HPCT with gating ties together in a natural way many previously
disparate phenomena. I had hoped to get it out before this, but other things intervene.

Martin

Date: Wed Jun 24, 1992 11:45 am PST
Subject: Re: Psychological research

[Martin Taylor 920624 14:10]
(Rick Marken 920624.1030)

Interesting that I should get Rick's response to my posting of this morning before I myself got my posting from the CSG-L distribution centre.

I'm not going to go up the statistics sidetrack here, since I think there is an underlying conceptual disagreement that won't get solved simply. But I want to answer one point.

On the basis of studies of control that can be reduced to tracking, studies that give correlations of 0.99+, Rick asserts that the near unity correlation can and should be obtained for all controlled percepts. I disagree, on three grounds, only one of which has to do with the possibility of measuring the controlled environmental variable. That first:

(1) Once you get to a reasonably high level of the hierarchy, it is quite possible that the experimenter has no ECS that controls exactly the same complex environmental variable (CEV) that is being controlled by the subject. If so, the experimenter cannot disturb the CEV in a known way, and cannot precisely assess the subject's control. The subject isn't controlling what the experimenter is disturbing, only something (possibly weakly) correlated with it. The experiment will show some, but not strong, correlation. Rick says, in effect, "fine. throw away the study and start again with a test on a new variable." That's OK for a solo scientist, but there is no way for that scientist to discover (i.e. build an ECS that controls) the CEV that the subject is controlling, other than by blind chance. Publication of the study provides some data for other scientists, who may be better able to control (disturb) the real variable. The study is not a blind alley, but a probe into the hills where some gold is found, if not the motherlode.

(2) For any CEV being controlled by the subject, there are probably other CEVs not the object of experiment but also being controlled by the subject and being disturbed by the environment. These other CEVs will, in all probability, induce conflicts in the hierarchy, which show up as noise in the control function. The experiment will show reduced correlation.

(3) Directly applicable to the VOT study: when a controlled variable is categorical, a wide range of values of the percept belong to the same category. If membership of the category is the controlled variable, then the physical values perceived will have a considerable variance, but will probably not approach the category boundary too closely. VOT is a highly distinctive feature categorically discriminating voiced from unvoiced consonants, but so long as VOT is not too short for an unvoiced consonant, or too long for a voiced one, the actual value doesn't matter much, and probably is controlled to indicate such things as mood and emphasis (which very probably changed during the course of the experiment). The gain function for the ECS is effectively flat (zero gain) for errors that leave the physical variable within the reference category, but steep as the physical variable approaches the category boundary. That leads to an increase of variance for the controlled variable if it is observed at a physical level (VOT = x msec), but not if it is observed at a categorical level (this sounds like a /p/, not a /b/). The experimenter cannot observe the categorical level, and hence sees high variance at the physical level, which is NOT an indication of lack

of control.

Comment on (3): Flat, categorical, gain functions with steep boundaries provide an escape from conflict in the hierarchy. So long as the induced variation in the physical variable does not go beyond the category boundaries, other ECSs can control it freely. The category-controlling ECS puts limits on the physical variations induced by these other, independent, ECSs, but it does not conflict otherwise with their control. Such flat controllers strengthen and liberalize the entire control hierarchy. Their central "don't care" region gives free rein for modulating functions, which would normally not be the object of experiments, and would therefore not be accounted for in assessing what everyday psychologists call "sources of variance." At higher levels, I suspect that most gain functions have this kind of non-linear characteristic.

> It's not a matter of getting less statistical variablility -- it's
>a matter of getting NONE. We know this can be done because we know that
>control is not a statistical process (how much variability do you
>tolerate when you control the you balance) and we have done it (see my
>book and the studies by Tom Bourbon and Bill Powers)

We know that control can be exercised only to the extent that information is available to the perceptual function of an ECS. And THAT is inherently a statistical process, so we know that control IS a statistical process. The only thing at issue is the relative amount of inherent statistical variability in the process compared to the range of the controlled variable.

Martin

Date: Wed Jun 24, 1992 1:06 pm PST
Subject: Misc..

[From Bill Powers (920624.1200)]

Martin Taylor (920623.2100) --

>One problem with the study is that the "adapting" VOT for the
>(synthetic) /r@pi/ was 50 msec, and the natural (after /i/) VOT for the
>subjects' /p/ averaged 57 msec, 7 of the 20 subjects being shorter than
>50 msec. So there wasn't much perceptual adaptation to be expected.

This is a coffee ground, not a peanut. If the authors reported that

> if the subject had listened to /r@pi/ rather than /i/, the voice onset time was shortened, particularly for /t/.

then they concealed the truth. This finding goes into the literature as a conclusion about what "people" do, and others will pick up this "fact" and use it as if it's true of everyone, and waste the time of Rick Marken, Bruce, you, and everyone else who feels the urge to explain phenomena. It isn't a phenomenon. Any explanation that purports to show why this "perceptual adaptation" occurs in "people" also has to explain why, in 35 percent of the cases, the OPPOSITE "adaptation" was observed.

I don't think that anyone exploring behavior using the HPCT model has any obligation to explain facts that aren't even facts. I think that from now on we should qualify all phenomena put up for explanation as to their status: are they false of any large fraction of the people in the study, or are they real robust phenomena that can reliably be demonstrated in essentially every subject?

Thank you for bringing this up. This is going to save me a lot of time. From now on, the first question I'm going to ask (if I remember to do so -- if not, remind me) is whether the phenomenon to be explained is actually a phenomenon. If it's not, I'm going to turn my attention to more profitable questions. From here on down, it's all peanuts.

Martin Taylor (920623.2115) --

>As soon as you allow same-level loops, you are into the realm in which
>oscillations and possible chaos lurk.

I concur. I concluded some time ago that there really aren't any "configurations of configurations" and so on. That's just a device of speech, one of the penalties for having to express parallel processes in sequential form.

I think we perceive legs, arms, seat, back and chair all in parallel. We can perceive any of these without the others. The Gestalt psychologists knew this 60 or 70 years ago. In truly hierarchically-related perceptions, there's NO WAY to perceive a higher-level thing without presence of the lower-level components, save in imagination. If there are no sensations, there are no configurations. This is an absolute requirement in the concept of hierarchical perception.

Rick Marken (920623.1830) --

Bruce Nevin asked

>Is adaptation a disturbance?

You replied:

>I call it a disturbance because if affects the perception of the
>controlled variable in such a way that action is required to maintain
>that variable at a reference level.

This is an important point. We tend to think of disturbances only as external forces that create error. But it isn't necessary for them to create error for the control system to oppose their effects. If you're holding a weight up at arm's length against gravity, you're controlling a perception of position. Now if someone else comes along and pushes upward on the weight, the error that's maintaining the muscle forces will decrease, and the muscle forces will also necessarily decrease. This decrease in your upward force opposes the upward force that the other person is adding. So the disturbance is opposed EVEN THOUGH THE OTHER PERSON IS HELPING YOU HOLD THE WEIGHT UP.

"Adaptation," which I agree is not a very useful term, occurs when the way a system acts changes. If you hold a dumbbell up at arm's length every day for half an hour, your arm muscles will "adapt" by getting stronger -- that is, more muscle fibers will be added. While you were weak, getting the muscles that existed to hold the weight up required producing a large driving signal from the spinal motor neurons. If the same driving signals were produced after the adaptation, the weight would not only be held up, it would be tossed backward over your shoulder. To keep the same "holding-up" occurring as before, your peripheral control systems reduce the output from the motor neurons; the same output now produces more force, so less output is needed to produce the force needed to counteract the weight of the same dumbbell, which is always the same force. Therefore the control system has decreased its neural output to compensate for the disturbance caused by the increase in signal-to-force transduction in the muscle.

A lot of what's called "adaptation" is just control. And a lot isn't. So the word isn't really of much use in HPCT.

Thanks, Rick, for anticipating my complaint about junk data -- I wrote my comment before getting to yours.

Hank Folson (920623) --

>1. How will a person (living control system) control when s/he does not
>know that s/he is a living control system, and does not know that
>everyone else is an independent control system that cannot be
>controlled without force?

Exactly the way he or she does now. People will try to make other people act to suit the actor's reference levels. The others will resist. The actors will try harder. The others will resist harder. Before long there will be blood all over everything.

>4. How will a person (living control system) control when s/he knows
>that s/he is a living control system, and knows that everyone else is
>an independent control system that cannot be controlled without force
>in a control theory _aware_ world?

The futility of (or the penalty for) trying to control other control systems of equal capability and intelligence will be obvious, and people will stop trying to do it, except when they have to (to survive, for example).

>Wouldn't the controlling actions be different?

Not so much the actions as the goals.

>1A. (A small but important sub-category) How will a person (living
>control system) control when s/he does not know that s/he is a living
>>control system, and does not know that everyone else is an independent
>control system that cannot be controlled without force, but has through
>their life experience developed a Systems Concept and controlling
>techniques that are compatible with PCT?

All methods of control are compatible with PCT. PCT doesn't teach methods of control; it explains the methods that already exist, and have always existed. There is no "technique" of control other than comparing your perceptions with your reference signals and converting the difference into action (via lower-order goals), as best you can. Some people have learned how to fool others into submitting to control, or taking advantage of their ignorance or beliefs to the same end. Control theory doesn't teach anyone how to do that. It explains what they're doing. It also explains why trying to control other people creates all the problems it creates, even when it seems to succeed.

>2. How will a person who has lived most of their life and had most of
>their experience in life in a Category 1 world as a Category 1 person,
>but then founded/pioneered/studied/was seduced by PCT, attempt to
>understand and control in our Category 1 world? Will s/he only use the
>new paradigm, or will s/he drift back and forth between paradigms?

We could take a poll. I drift back and forth between paradigms, although as the years have gone by (more for me than most others), I've spent less time in the system concepts I was raised with and more in a new one. I think that when it comes to the crunch, I now opt for the new one almost all of the time. This will be much easier for people who don't have to overcome the old ways to appreciate the new.

>3. How will a person who has been raised in a Category 1 world, and
>whose Systems Concept is rooted in Category 1, and who may even have a
>vested interest in the world remaining a Category 1 world, control
>(still in determined ignorance of their being a control system) when
>faced with the PCT concept of the world?

Stick around and watch. Most of them will probably grow old and die without ever changing their minds, while continuing to act as perfectly good control systems without realizing it. They will exert great efforts to protect their system concepts against disturbance. They will try to attack the disturbances at the source, or to prevent them from being effective, or to bring greater forces to bear and overcome the disturbances.

>Would you gain any new insights by using these categories were you
>reread and analyse again what is going on in the recent posts on the
>newspaper organization, leadership, Chimps, the Jury comments?

Probably. But don't be shy: tell us what new insights you think we might get.

RE my quote: don't overlook the part that says "... if we want to find the significance of the first new concept of human nature since Descartes." I'm trying to say what we need to do IF that's the goal. For a lot of people, it still isn't.

Bruce Nevin (920623.1234 --

An interesting study on the fish. How would you reframe it in PCT terms? For starters, I'd recommend considering what "aggression" means as a control phenomenon, and how the relationship of "domination" might work.

There's an interesting parallel here with some vague ideas I've had about cell specialization during maturation. All cells have all genes, apparently. But when a certain number of cells with a particular gene become active, the result, according to the conventional interpretation, is that additional cells with the same genes are turned off, so we get livers, fingernails, gonads, and so on.

But suppose that these genes represent control systems controlling for some aspect of the common milieu. When there are enough control systems for a given variable, say an enzyme, the controlling systems will experience zero error, near enough. This means that all the cells with a reference level for that variable that's less than the amount being maintained will experience truly zero, or even negative, error and their outputs will turn completely off. So there's no need for a third party to actively switch genes on and off: the controlled variable itself will do that.

Suppose that all the male fish have an intrinsic reference signal that says, in some equivalent form, "What we need around here is a flashy horny male to keep this society in line." The actual controlled variables would be more specific things that have this situation as a consequence. One would suppose, then, that the male who maintains these things at the highest reference level would satisfy ALL the males' reference levels for those things, and their control systems would atrophy through disuse. But when that obliging male dies, suddenly all the males begin to experience this intrinsic error, and reorganize until one of them manages to satisfy all of them again. "Reorganize" may not be the word -- this seems too regular a relationship for reorganization. But you get the idea.

The guy with the highest reference level, and I suppose the highest loop gain, ends up doing all the work.

Best to all, Bill P.

Date: Wed Jun 24, 1992 2:17 pm PST
Subject: Re: what's the difference?, psych research

[From Rick Marken (920624.1320)]

Obviously, this is a lot more fun than working.

Martin Taylor (920624 13:00) says:

>Oh, wow...disinformation piled on misinformation!

Oops. Hit a controlled variable.

> I know that many people
>think of aftereffects as the consequence of fatigue, but they can't be, at
>least in most cases and perhaps in all.

Just meant to describe 'em, not explain 'em. Let me guess; have you published in this field?

> It is NOT true that if you adapt to up movement and
> then show up it still looks up.

This was not intentional disinformation on my part. It was just based on my own experience. After staring at a waterfall it still looks like it's going down; after staring at green it still looks green. There may be adaptation effects but they are not phenomenally obvious to me. But I believe you if you say it happens.

> A naive view based on
>"adpaptation=fatigue" cannot work because it predicts a lot of phenomena
>wrongly. "Adaptation=improved perceptual precision" accounts for quite a
>few of the phenomena that "adaptation=fatigue" does not, and predicts
>numerically as well as qualitatively.

OK. If you say so.

> In the red-green case, it is possible (even likely) that fatigue plays a part.

Well, so my "fatigue" story wasn't all THAT bad, after all.

>In the movement case, it is less
>probable, and when we come to the more shape-based aftereffects, it is not
>likely at all.

I didn't mean to step on any theoretical toes. I don't really have any preferred explanation of perceptual adaptation. It's just fun to experience it.

>I don't follow Rick's analysis of the VOT detector system at all:

I was just thinking on my seat -- in terms of what Cooper/Nagy might have had in their head when they did this adaptation procedure. I was not proposing any explanation of my own. And I typically leave the "disinforming" to my friends in the CIA.

>If there really is a VOT detector, it asserts whatever VOT is appropriate
>for the phoneme in question as a reference.

No comprendo. How do detectors "assert" anything?

>Bruce is right. There are several different effects of adaptation.
>Adaptation is a procedure, not a "perceptual effect".

I was just describing the procedure and the phenomenology. I used the fatigue story to describe what seems to happen. Now I wish I'd never said anything about "fatigue". Geez. Mea culpa. Mea culpa.

>Rick often proclaims that all psychological studies done outside the
>control paradigm are worthless.

No. I said that analyzing random noise is worthless -- no matter what paradigm you use. The chances of getting random noise in the study of behavior, however, is almost guaranteed if you use the s-r paradigm.

> This shouldn't give him the right to assert
>his own view of the world that they study, in contradiction to the results

>they obtain. You can't claim better truth by throwing away data than by
>looking at what data you have, however wrongheaded the data gatherers might
>have been.

I did that? I was just saying that you will see a red square on white paper after you stare at a green square. Any endorsement of a theoretical model of perceptual adaptation, either implied or stated, was purely coincidental.

>Sorry to be so harsh, Rick, but that message really seemed wrongheaded, if
>not bull (headed).

It's OK. I hope your perception of what constitutes an appropriate way to explain perceptual adaptation is back under control.

Now, to create another disturbance:

Martin Taylor (920624 14:10) says:

>On the basis of studies of control that can be reduced to tracking, studies
>that give correlations of 0.99+, Rick asserts that the near unity correlation
>can and should be obtained for all controlled percepts.

Nope. I said (or meant to say) that the criterion for what constitutes a scientific fact in psychology should be far stricter than it is. I think a reasonable goal is correlations of .99+. This can be done when you are studying control -- at least when you are studying variables that can be quantified relatively easily. It should even be possible with higher order variables that are harder to quantify (David Goldstein and Dick Robertson did a study of control of "self esteem" where they got .99+ correlations). It can be done. It must be done if the study of living things is ever going to be a science instead of a dice game (not that there's anything wrong with dice games).

The three reasons you give for why one can't expect to get perfect correlations even when studying control are ok. But they have nothing to do with current research in psychology. The goal of research should be high quality data -- always. Nearly all research in psychology provides low quality data. The fact that this data is collected in research that is done from the wrong perspective is irrelevant to my point -- which is that there is precious little to be learned from looking at noise. The JASA VOT study illustrates this point. Now, you can come up with all kinds of reasons why they couldn't have gotten better data -- or you can just go out there and get good data. I say "go with the second option"; the first option does nothing but try to justify trying to make sense of worthless garbage -- salt or no.

>We know that control can be exercised only to the extent that >information is available to the perceptual function of an ECS. And THAT >is inherently a statistical process,

What is "THAT"? Information? The perceptual function? What is it about control that is "inherently statistical"?

>so we know that control IS a statistical process.

When a non-statistical model accounts for 99.87% of the variance in the variables involved in control I think it's fair to say that the "inherently statistical" part of the process of control is not worth losing much sleep over.

Go for the QUALITY data.

Regards Phaedrus

Date: Wed Jun 24, 1992 3:49 pm PST
Subject: Good data, bad data

[From Bill Powers (920624.1500)]

Martin Taylor and Rick Marken and Bill Powers (920624) --

RE: quality of data

Yeah, let's pick at this scab a little more, but let's try to minimize the bleeding. Some bleeding, I think, is unavoidable.

Martin Taylor (920624.1300) --

I did an experiment with aftereffects of screens full of moving dots, with Pat Alfano (last year). Same principle as yours: the subject moved a control handle, immediately after the adaptation period, to make the dots appear to stand still for a few minutes. We got beautiful curves showing the decay of the illusion for lateral, vertical, rotational, and convergent/divergent movement, showing quantitative effects of speed before adaptation. I just matched an exponential to the data -- no analysis as fancy as yours. The point was not so much to do a fine analysis as to look for differences between people with and without motion disorders.

Just as Pat was about to start her PhD thesis using this experiment, BOTH SHE AND I UTTERLY CEASED TO HAVE ANY MOTION ILLUSIONS IN THIS EXPERIMENT. The offset of the illusions occurred about a week apart for us, mine disappearing first. As far as I can tell, nothing about the apparatus changed, although I beat my brains out trying to find some difference. Pat, incidentally, came out of this essentially cured of the debilitating motion disorder that had prompted her to do the study in the first place!

I really think we both ended up reorganizing, through too many hours of familiarity with the experimental situation. So your warnings about the hazards of experimentation fall on receptive ears, here. Pat, by the way, had to go to a completely different approach, poor thing. Graduate students should stay away from control theorists.

Well, on to the real subject.

Here's my main objection to statistically defined phenomena, in a nutshell. If someone comes up with an effect and I decide to try to model it, the model I produce will embody the mechanism by which I propose that this effect could be produced. So by its nature, the model predicts the behavior of ALL subjects in a given experiment.

If, however, I then find that in fact, only a majority of subjects show this effect, while one or two or twenty show some quite different behavior under the same conditions, where does this leave the model? Now the model, instead of fitting the data, predicts clearly wrongly for a substantial proportion of the subjects. The model, in other words, has failed. Fixing the model to fit all the data might require only a small change -- or more likely, it may require completely abandoning the basic premise and starting over. Even one clear counterexample is enough to do in a model, although if there's really only one, and it can't be reproduced, I'll admit that I might hang onto the model a little longer.

If, however, a model CONSISTENTLY fails to account for some even small number of observations that it's supposed to predict, there's no choice but to track down the reason and modify the model accordingly, or scrap the model and start again from a different set of assumptions. If the data themselves are irreproducible to some extent, then modeling is futile.

In your 10:30 post you say

>Once you get to a reasonably high level of the hierarchy, it is quite >possible that the experimenter has no ECS that controls exactly the >same complex environmental variable (CEV) that is being controlled by >the subject. If so, the experimenter cannot disturb the CEV in a known >way, and cannot precisely assess the subject's control.

This may be true, although until we actually do PCT experiments it may be premature to borrow trouble. My question is what one does about such a situation when it's encountered. Your answer is one I've heard many times before:

>Publication of the study provides some data for other scientists, who >may be better able to control (disturb) the real variable. The study >is not a blind alley, but a probe into the hills where some gold is >found, if not the motherlode.

I think this is a highly idealized version of science. If scientists in the behavioral sciences actually did take equivocal results from other people, replicate their experiments, and eliminate sources of variability bit by bit, this might be a reasonable approach. But in fact that's not what happens. Replication is all but unheard-of -- everyone wants to do his or her own jazzy experiment, not slog along cleaning up behind someone else a la Bullwinkle. Even when replications are published, there's always some critical change in conditions, methods, subjects, or something -- the temptation to improve on the original is apparently irresistible. And as to trying to reduce the unaccountable variability -- well, have you EVER seen a study like that? Have you ever seen a study in which somebody said "Gee, X's experiment left 20 percent of the variability unaccounted for. So I tried to find out where it came from, and now only 10 percent is unaccounted for." Maybe you have; you read more of this stuff than I do. But my strong impression is that once a study has been published, with its findings pronounced statistically significant, all variability vanishes from the description of the phenomenon, and from then on, as far as the rest of the scientific world is concerned, the phenomenon has been established as a fact true of all subjects who fit the population description. The more interesting and striking the result, the less it matters that large numbers of people don't fit the description. Tell me this isn't true, if you dare.

That's one objection: elevation of preponderances to universals.

Another objection I have is simply that the levels of correlation accepted in the literature (Gary Cziko has lots of data on this) are abominably low. Rick cited a study on VOR in which r = .34 was the basis for reaching a conclusion about how "people" speak. Yet if the AVERAGE VOR was taken as the predicted VOR for all people in the group regardless of treatment, prediction of an individual's behavior would be sqrt(1 - r^2) or 94% as accurate as would be a prediction based on the supposed effect. I simply refuse to accept this sort of thing as data. I'm supposed to come up with a general model of behavior that makes predictions that are 94% useless? This whole VOR business could be the result of the way people move their mouths from one configuration to another, and differences in this manner of producing output could have no importance at all.

Rick Marken said this well, but I'll say it, too. The behavioral sciences have simply grown into the habit of accepting data that are completely inadequate for science. They've blamed variability on their subjects instead of blaming a wrong concept of what's happening.

I know that it's always possible to make excuses for bad data, even in control theory. You can say "Well, maybe the subject was controlling for some other things that interfered with the experiment." This may indeed be true -- but you've offered an hypothesis, and now it's up to you to show that this is in fact why the data were bad. WHAT other things was the subject controlling for? Do the experiment, show that this is in fact the right explanation. And then fix the model so it explains, instead of making excuses for it. I'm not going to accept even an impeccably phrased HPCT excuse for poor predictions as a reason for keeping a model of a specific behavior. I say go away and work on your model until it DOES predict correctly. And for heaven's sake, don't publish until it does. Why clutter up the literature with junk before you can convince others that you have it right? All this stuff about leaving hints for others to follow up sounds very nice, but the most likely fate of such work is to be forgotten the instant after appears in print -- if it's ever read at all.

>(2) For any CEV being controlled by the subject, there are probably >other CEVs not the object of experiment but also being controlled by >the subject and being disturbed by the environment. These other CEVs >will, in all probability, induce conflicts in the hierarchy, which show >up as noise in the control function. The experiment will show reduced >correlation.

You see? Even you do it. What other CEVs? Find them, test them, and get rid of the variability. Don't just accept a "reduced correlation" as the inevitable consequence of things we can't help. If there's a reduced correlation, you can't use the hypothesis as a fact in any grown-up scientific argument: bad data is bad data no matter how much it's not your fault. If an hypothesis is not very likely to be true in any given test of it, it won't get any more true by being used in a deductive argument.

As I said, I think you're borrowing trouble. When you think up a sound PCT experiment, you're not going to find a lot of interfering variables reducing your correlations to where you have to explain why they're so low. When you hit on the right way of doing the experiment, you're going to get good data. I mean you, generic. If no matter how you try, the data still come out bad, then you've got the wrong idea or you're into something more complex than you can deal with at our present stage of understanding. Do you think Galileo was ready to explain how compasses work? There are some things that just have to wait a while until we build up a base of solid knowledge.

My attitude is this: let's explain what we can explain, and not lower our standards just to appear wise about things we don't understand yet. As Rick said, it's all right to say "I don't know."

Let's explain simple aspects of behavior with high precision. In that way, we will leave behind something on which others CAN and WILL build. The longer we or our descendants stick to this principle, the greater the cumulative effect will be and the more complex will be the behaviors we can confidently and accurately explain. The real sin in the behavioral sciences has been the pretense of knowing what nobody actually knows yet. Go ye and sin no more.

Best, Bill P.

Date: Wed Jun 24, 1992 4:06 pm PST
Subject: causation

Restating my initial ideas:
  Outcomes are influenced by the environment and reference signals.
    Outcomes are determined by reference signals.
    Actions are influenced by the environment and reference signals.
    Actions are determined by the environment.
    Actions control outcomes.

First, yes, indeed I made these statements with the assumption of low intrinsic error.

Second, these statements were only meant to describe relationships within one isolated control system, not considering the hierarchy as a whole. I intended on addressing that issue next, but Bill already beat me to it. I'll say more on that later.

Third, the terms I use (influence, determine, control) are defined by Bill (910509) as follows:

"A influences B if A is one of the several variables on which the state of B depends...A determines B if, given A, B is completely predictible...A controls B if, for every disturbance applied to, A changes its influence on B in such a way as to counteract the effect of the disturbance on B."

Bill and I either disagree or are not using these terms in the same manner. For instance, Bill states (920617):

"(a) The reference signal and input function determine the state of the input quantity in the environment...(b) The reference signal and external disturbances jointly determine the output quantity....(c) The total disturbance, composed of d1a and d1b, and the second-level reference signal, jointly determine the second-level output, which translates into qi1b in this case."

Here, I would use "influence" where Bill uses "determine." In (a) I would use "influence," not "determine" since the input quantity is predictible solely by knowing the reference signal. In (b), "influence" is appropriate since the output quantity is not predictible on the basis of the reference signal. In (c), I would save the word "determine" for the reference signal.

If we are using different terms with the same meaning in mind, then we agree. Do we agree, Bill? Or do you mean "determine" the way I (you) define it above? Perhaps the problem resides in tht I have difficulty making a distinction between an input quantity and input function, an output quantity and output function, a disturbance quantity and disturbance function.

Now, back to my second paragraph. Bill states (920617)

"(b) The reference signal and external disturbances jointly determine the output quantity.

Statement (b) says that given a constant reference signal, variations in the disturbance call forth specific variations of the output quantity or action, in the manner of an apparent causal relationship. The determining effect of the disturbance on the output, however, is subject to the condition that the sum of disturbance and output effects always equal a particular value: the value of the input quantity determined by the reference signal. This balance point, therefore, can change if the reference signal changes. This is why the action of the system is JOINTLY determined by disturbances and the reference signal, and not exclusively determined by either."

I would further support my statement that the output is determined exclusively by the environment (inluenced by the environment and reference signal) and NOT jointly determined because the change of reference signals Bill refers to is the result to higher level systems changing to counter environmental disturbances. At some level the reference signal is stable and the changing reference signal (level n)/output (level n+1) are determined entirely by the environment. The changing reference signal which Bil says jointly determines the output is itself determined by the environment, albeit an "internal environment" depending on the level of analysis." This seems consistent with the conclusions arising from the dialogue between Bill and Greg a month ago on Autonomy.

I agree with Bill when he states:

"At any level of interpretation, statements (a) and (b) will hold true -- but with many systems at each level, each level has to be considered anew. When a single control system at one level receives reference signals from several higher systems, there can be no simple relationship between disturbances of a given higher-level perception and the resulting change in the lower-level net reference signal."

and in some sense I agree with Bill when he states (910615) that perhaps ordinary language limits understanding of such concepts. On the other hand, it is not so much ones language that makes understanding difficult but the dynamic nature of the beast itself. Words or no words, volitions or no volitions, the interactions within complex dynamical systems are (to be understated) difficult to understand. I can at least hope to put words to what I do understand, so I disagree that such attempts are not useful--it at least allows us to talk about it.

Finally, Bill states:

"Both of these statements describe apparent causal relationships, which are different from those that actually exist in the control system. That is, the "determination" takes place through a path different from the one that appears to exist. These two statements describe appearances, but not the actual organization of a control system."

I have tried to make a distinction between "actual causal relationships" and "instrumental (appearances) causal relationships" and have concluded for the time being thta any distinction made between the two is an instrumental one, not actual. The issue still perplexes me, but my guess is that the above is the case because the concept of "cause" is useful for organisms such as ourselves in adapting to our environment. It may be akin to making a distinction between "actual color" and "appeared color" (excuse the awkward wording) (This is probably a poor analogy though). If I inject you with a drug, do I say that the cause of death was injecting the drug, or any one of the number of chemical reactions which occured along the way to the point of death? I don't think any answer is better than another. Some causes are more remote/proximal than others, but neither is more the cause than the other. I am justified is saying that _a_ cause of the grass dying was the application of fertilizer even if a more proximal cause was some "chemical reaction x" of which I have no knowledge. so when Bruce states (920616): "How do reference signals influence outcomes other than by way of actions?" I agreee with his point but still contend that I need not mention the entire causal pathway. I think "cause" is closer to "influence" than "determine" for related reasons.

So, I still contend thta my original formulation is accurate for a single control system, using Bill's definitions (910509). I like what Bill said (920617):

"Considering only the first-level system, we still have the reference signal determining the input quantity, now qi1b. This means that the output of the second-level system is, as far as second-level control is concerned, not qo1 but qi1b. The input quantity of the first-level system, not the output quantity, will appear to be the action of the second-level system."

I would restate that as "action (level 2) = outcome (level 1)" which can be used to derive the relationships within the hierarchy. This may become extremely complex, and perhaps too linear, but still useful nonetheless.

Oh, one last thing in response to Bill's post (920615):

" A lever has the property that pushing down on one end makes the other end go up. So it determines that IF you want the far end to go up, THEN you push down on the near end. Skinner used the term "contingency"

for this: movement of the far end is contingent on movement of the near end. Not a bad term.

But the lever neither causes nor influences anything by just sitting there.

The variable position of its far end, given the lever's properties, is influenced in a particular way by the action of moving the near end. It is also influenced by independent forces acting anywhere along the lever. The lever doesn't determine whether any such actions or forces will occur."

I do not catch your point in your example with the lever, but I will say this: it is not correct to say that the variable position of the far end is influenced by moving of the near end. Causal relationships must be temporal. An increase in air temperature within a balloon does not cause an increase in pressure (if volume and n remain the same). They are interchangable since PV=nrt. No temporal relationship exists. Pushing down on one end of the lever is identical to making the other end go up; it does not cause it. Pushing causes "going down"/"going up"; it is not the case the pushing causes "going down" which in turn causes "going up."

Carpe' Diem

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: Wed Jun 24, 1992 8:19 pm PST
Subject: Re: Chimps; Xs; HPCT as tool

(sibun 920624)

>Second, the description of the social interactions employs stereotypes >that probably don't reveal what even human participants are controlling >for in such situations. There's a bit of cynicism in the description: >I'll bet Carl Sagan never stood in a crowd and reached out to touch a >charismatic leader like George Bush, or that he ever bowed down in awe >before any other person. That's just his (their) stereotype of how the >common masses behave. If S&D could get inside the heads of most "common" >participants in rituals and ceremonies, they might be shocked at finding >how much of the same sort of cynicism is there. If Sagan could hear >people reacting to him when he strikes an imposing pose on television, >he might be furious, as well.

nevermind the stereotype that it's all *male* actors in the scene....

--penni

Date: Thu Jun 25, 1992 9:34 am PST
Subject: HPCT research

[From Bill Powers (920625.0830)]

Some general thoughts on the VOT etc. controversy and HPCT research in general.

There's a problem that hasn't come up for discussion yet. Suppose we say, when a person reacts sharply to a criticism by another person, he's controlling for self-esteem. Good, a PCT hypothesis. So we set up an experiment in which a person is asked to respond to a series of statements, some critical and some not, and we record the responses and evaluate then as correcting an error in self-esteem, or not. We do the experiment. We get a correlation of 0.2 between critical statements and defensive statements. What do we do then?

Obviously, this is a poor result. If we were to follow custom, we'd just go on to another hypothesis and hope for better luck. If the new hypothesis gave us 0.6, we'd become happier, and publish.

As I see it, however, the first experiment hasn't been finished. What's missing is any attempt to improve a model. There is a model: it says that self-esteem is a perceptual variable that the person is controlling for. A critical statement is a disturbance of that variable; a defensive statement is an attempt to correct the disturbance. So the model says that disturbances should elicit opposing efforts; critical statements should elicit defensive statements. If we input all the critical statements as disturbances of the model's controlled variable, the model will produce defensive statements. But the real subjects behave in a different way, given those same critical statements, the same disturbances.

It wouldn't take long to see some possible mistakes in the model. We're assuming that everyone wants high self-esteem. We're assuming that certain statements critical of a person would be perceived as disturbances of self-esteem (instead of simply indicating a hostile experimenter). We're assuming that the "critical" statement implies that the person is worse than the person already thinks he is (maybe the person would be relieved at being called lackadaiscal if he think's he's incompetent). We're assuming that the critical statements succeed in altering self-esteem. And we're interpreting the person's responses as defensive or non-defensive without knowing that person's perception of his or her own responses. In other words, most of the parameters of the model are being filled in with imaginary data. No wonder.

By examining what the person actually did, and comparing it with what the model did, we can begin to see how to alter the experiment and the model. At the least, we ought to try to find out what certain statements mean to each subject -- whether, in fact, saying "You make too many mistakes", if this statement were true, would worsen rather than improve this person's self-esteem. We should find out whether a response like "Yes, you could be right" actually indicates agreement, or conceals a resentment that the person is ashamed of. We should try to find out if the person has high self-esteem already, and wants to maintain it, or has low self-esteem which is simply confirmed by criticism.

So the poor results in the experiment should tell us that the model is inadequate, but shouldn't discourage us from trying to improve it. As we think of reasons for the failure of the model, we can change our assumptions, sharpen up the interpretations of questions and answers, make the meanings more relevant to each subject, and so on.

This is a lot of work and will take a long time. We may have to modify the model and repeat the experiment from scratch 20 times. Is this worth doing?

If we anticipated that going through all this labor and spending a couple of years on this question would raise the initial correlation from 0.2 to 0.6, and $p < 0.055$, the conclusion might be that it's not worthwhile. Of course this means that we didn't think the original question, do people defend themselves against criticism because of disturbances of self-esteem, very important in the first place -- if we're willing to give up after the first try, we couldn't have been very interested in the results, anyway. I suspect that this tends to be true in graduate schools where the point is to get the degree, not necessarily to discover something important.

But suppose all our experience told us that by iteratively experimenting and modifying the model we should eventually come up with a correlation of something like 0.98. Think how this would alter our attitude toward the initial correlation of 0.2, or even 0.6. We would look at that abysmal result and think "Oh, crap, that's awful! What did I do wrong? What kind of stupid mistake did I make?" We'd be shocked into taking a fresh look at every assumption, every detail, looking for the hole in our reasoning. We wouldn't even THINK of publishing.

When you're used to seeing correlations of 0.6, getting such a correlation doesn't discourage you -- it pleases you. That's what you expected to get. You think you've learned something; you think your hypothesis must have a grain of truth in it. If all the people you work with and all those who read your work and judge it have the same kind of experience, they will accept your work and congratulate you for it.

But if you're used to seeing correlations of 0.95 and up, and your colleagues and critics are too, the same results look entirely different and your subsequent actions become entirely different. You think there must be something terribly wrong with your hypothesis, or your interpretations, or your methods. You roll up your sleeves and get back to work. You will work all the harder and longer if your experiment concerns some aspect of human nature that you think is very important to know the truth about. You don't ask "what is truth?" and you don't say "all you can really prove are negatives." Those are the things that people say when they expect correlations of 0.6 and are uneasily looking for an excuse. You expect to know the truth, plus or minus three percent. This isn't to say that another truth wouldn't work better. But you know the difference between true and false.

Suppose that you're a psychologist just starting in with HPCT. You hear a lot of bragging: "We can get correlations of 0.997 that hold up with predictions over a span of a year." Or "When you do a real PCT experiment, you get an exact match between the model and the real behavior." Intrigued, you ask to see some of these wonderful experiments and models. And what you are shown are some people sitting in front of a computer screen wiggling control sticks.

Oh. Is that what you mean? Well, er, um, that's very nice. I'm sure this is very interesting, predicting how people wiggle control sticks. Keep up the good work. See you later.

Thus is lost the meaning of a correlation of 0.997. The psychologist looks at what we would like to have explained -- the causes of anxiety, the causes of social friction, the problems of living in general -- and compares it with what PCT explains so well, and decides that PCT isn't concerned with serious subjects.

What is hard to get across is that PCT research is truly starting all over, from zero. It's concerned with establishing some facts about behavior that are as accurately known as most facts in physics -- but not just by doing physics. There aren't any such facts in the whole of the psychological sciences outside psychophysics. There is, therefore, no such thing as psychological science (on the same plane with physics), and there can't be until the base is constructed. PCT research is constructing a base at the same level of detail that Galileo explored by rolling balls down inclined planes and timing pendulums with his pulse. Galileo might well have bragged that he could now predict physical phenomena with unheard-of accuracy, over long spans of time, with nary a failure. And people interested in building bridges, fortifications, projectile launchers, and ships might well have investigated his claims, seen what he was actually doing, and said "Very interesting, I'm sure. Keep up the good work. See you later."

When you've thought up an experiment to test a model, carried it out, and found a correlation of 0.997 between what the model does and what a real person does, there's only one response: jubiliation. You have actually discovered a real true fact of nature, a high-quality fact, and fact that sticks up out of the mass of other facts like a lighthouse. You have found this through a direct confrontation with nature, in which nature could and did behave any way it pleased. And it behaved exactly the way you thought it would and predicted that it would. This is very heady stuff and there's no experience on earth like it. This has to be how Galileo felt.

The last thing in the world you worry about is that the behavior you've predicted isn't very complicated. Nobody has EVER predicted ANY behavior, even the simplest sort, with this kind of precision. Now you know something that nobody else has known: it's possible to do this. At this moment, every experimental result in the behavioral sciences that you've ever heard of changes drastically for the worse. Fact? You call that a FACT?

If you can get 0.997 in a simple experiment, maybe you can get the same result with a slightly more complicated one. Yes, you can, it turns out. You can even have two people controlling the same display, with two interacting models predicting their behavior, and still get 0.997 correlations. What about four people? What about having them control something a little more complex, with somewhat more complex actions? Yes, and yes. The correlations hold up. The model continues to work. The facts stick up like lighthouses here and there, in a sea of low-quality facts. They're not skyscrapers yet. But they show the way.

Once you've set foot on this road, you can see that it leads where we want to go. Eventually it will lead to a solid reliable understanding of all that is possible to know about human behavior. There's no point in trying to skip ahead and guess how it will all come out. There's no point in using methods that produce bad data and bad guesses and lead to knowledge that has only a minute chance of being correct. Certainly, those bigger problems are important. Certainly we need to solve them, as soon as we can. Certainly, we have to go on trying to cope with them using experience as a quide where we have no understanding. But if it's a science of life we want, somebody has to aim for 0.997 or better, and keep aiming for it no matter how slow the progress. Because only in that way are we going to understand and not just fool ourselves into believing that we understand.

Only a handful of people in the CSG has had the experience of making and testing PCT models that work. For the rest, it's been a vicarious experience. A lot of people have taken up the PCT ideas and have tried to apply them creatively to understanding behavior, even complex behavior. But everyone needs this basic experience of what it's like to make an essentially perfect prediction. When you have had that experience, you're no longer satisfied to conjecture about what people MIGHT be controlling for. You get the notion that it's possible to find out what they ARE controlling for. This movement isn't going to take off and fly until a lot of people start getting that idea: that not only can you guess, you can really find out what's going on. And until they start doing it.

Best, Bill P.

Date: Thu Jun 25, 1992 9:38 am PST
Subject: Chimps

[From Bill Powers (920625.1100)]
RE: chimps

Penni Sibun (920624) --

>nevermind the stereotype that it's all *male* actors in the scene....

Don't fret, Penni. Maybe it's true that in chimpanzee societies, male stereotypes exist. What more can you expect of chimps?

Bill P.

Date: Thu Jun 25, 1992 11:24 am PST
Subject: more stat rap

[from Joel Judd]

Martin, Rick, Bill, etc.

Here's my take on the stat thing having looked through literature in my field (SLA).

I think Bill gives too much credit to editors and other "mentors." Replications are not published becasue thay are not wanted, in many instances. Graduate training includes spoken or unspoken directives to finally "go where no man has gone before." A replication does not provide committees with evidence that a student is original and inventive, it shows an advisor's lack of progressive influence on the advisee.

But can a replication (of a study using group statistical data) even be carried out? Think about it. Taylor (1958) and Runkel (1990) have. I look at the studies in SLA, and I don't know if even 15% of them could truly be replicated. No one has true random samples. This is the *basis* for methods of relative frequency. Nobody has them!! How can I replicate a non-random sample? Maybe it's actually easier--I can use anyone I want! Runkel goes even further: Suppose I want to replicate a random sample of a population. Can I repeat a study carried out at time X six months later? Six years later? Can I claim anything for a true random sample when I take eight months to complete the experiment--haven't the population characteristics changed by then? But these are moot points for me, I repeat, no one (of widespread notoriety) uses random samples in studying L2 acquisition.

The point of studied examinations of social science research and group statistics (like the two references listed above) is: you're using the wrong tool to find out about individual functioning. Relative frequencies were never intended to be used to uncover explanations (i.e. theories) of human functioning. By their nature, they CANNOT DO IT. The facts (at least in my field) are that there is widespread looking the other way as the methods are misused and misinterpreted, that shoddy results become SLA fact over time, and that ever more powerful and hopelessly uninterpretable techniques are applied to language learners each year. These should be clues that something is amiss.

Date: Thu Jun 25, 1992 3:17 pm PST
Subject: aftereffect experiment

[From Pat Alfano]

Dick Robertson showed me your comments on the aftereffect experiment, and I wanted to raise some still unanswered questions. I have had a few people try the experiment, and have tried it on different computers, no one is getting any aftereffects. The first version on the program you made for me induced aftereffects in everyone who tried it. Could something have changed in the later version to eliminate the aftereffects? Would like to use the program to see if visual aftereffects are related to motion sickness. What do you think the problem could be?

By the way I recently received my degree from DePaul. Feels good to be a Ph.D. Would feel better if I had a job. How do you like Colorado? We re in the process of selling our house, will probably stay in Chicage. Hope you and Mary are well. Hope to 'talk' to you again.

Pat Alfano

Date: Thu Jun 25, 1992 3:55 pm PST
From: Control Systems Group Network

[from Gary Cziko 920625.1625]

Last weekend I was at Lehigh University in Bethlehem, PA to consult with Donald T. Campbell about my book (which is very Campbellian).

While there, I was able to meet with Mark H. Bickhard, Luce Professor of Psychology, Philsophy, Robotics, and Counseling. Bickhard has written extensively on what he calls "interactivism" and has a view of cognition and behavior which I find in many ways consistent with HPCT. But he finds that HPCT is not adequate to deal with all important aspects of learning and development.

I am sharing with CSGnet a transcript of that part of our conversation dealing with HPCT. I think CSGnetters will find his comments of interest. Bickhard seems to be one critic of HPCT who understands the basics of what HPCT is all about.

While Bickhard will probably not join discussions on the net, I may send him reactions that I feel he may be interested in. Perhaps this is one way to coax him to share more of his ideas with us.--Gary

===
Mark Bickhard (MB) Interviewed by Gary Cziko (GC) 920621

MB: The knowing levels,I would want to claim, are levels of reflective consciousness.

Now within any given knowing level, there could be other principles of hierarchicalization but they won't get you to a new knowing level, like potentially servomechanism hierarchies. But the relationship between one level of a servomechanism hierarchy and another level is not a relationship of epistemic aboutness. It's not a reflectivity.

GC: How do you get that? Where do you get that reflectivity of that knowledge?

MB: Well, the basic notion is that if knowing is interactive then the sense
in which a system can know the world is in terms of interacting with it, or at least being competent to interact with it, would be also a sense in which some second-layer system could know the first one by interacting with it. And, in fact, there would be reasons why that would be adaptive because there are new things to be known at that level and there will be new things to be known at the second level but that will require a third level, and so on. But in order for that to happen, you've got to have a system that actually interacts with the next lower level as differentiated from a system that simply calls upon it in a control-flow sense, like from one servomechanism to another.

GC: In Powers's model you have sitting on the side a reorganization system and when there are errors this thing pushes a button that says "change in some way." Is that not an interaction of the type that you're talking about?

MB: Well, it's a very, very limited one. As a matter of fact, I have a model of the macro-evolution of this second layer system and the first limited version of it is in fact a learning system which is at the meta-level with respect to the system it operates on. But it's very computationally limited and is also a very interactively-limited system.

It can't do very much with the level below. I mean, what it does is terrific, but it can't do it a lot. The second one is an elaboration or modification of a learning system plan and I argue that it constitutes an emotional system and it's more powerful than just a learning system. And in the third one, I would argue, is a full reflective knowing system. And each one of those is, so I argue anyway, an increase in adaptability over the preceding step. Each one is a modification, not necessarily a trivial modification, but a modification of the preceding step, and in that sense I argue that they constitute a macro-evolutionary trajectory. And in fact if you look at evolution, that is in fact the order in which they seemed to have evolved.

GC: So the first one is a limited type of meta-system?

MB: Well, the very first one is a knowing system that doesn't learn much or
can't learn much. (yeah) And then comes a system where progressively more powerful learning capacities and at some point you get an emotion system and at some point you start getting reflectivity.

GC: Since we've sort of drifted into the hierarchy and talking about control to some extent, I'm really fascinated with Powers's model and I guess there's a number of reasons for that. But it's this notion of purpose which I find really intriguing and how these fairly simple servo-mechanisms seem to have purposes and resist disturbances. James talks about how organisms and people are able to obtain consistently certain things by varying behaviors, so his notion of controlling perception I find intriguing as a basic model. You talked a little bit last night about the problems you see with that and you mentioned the problem of correspondences--there is no mapping that can be used in this way that is going to be . . .

MB: I claimed that it faces ultimately a number of problems one of which is
the inverse of behaviorism. You can have very, very simple systems whose competencies are an infinite class of possible behaviors. So there's no way to characterize that system in terms of its behaviors as long as you're restricting yourself to finite characterizations. The only finite characterization that has been given that system is system organization, not behaviors. And exactly the same point holds for input. You can have very, very simple systems that can recognize, detect, differentiate (or whatever kind of word you want) infinite classes of inputs, and so in terms of input, there is no finite characterization of that system possible. That's point one.

GC: I have difficulty with that.

MB: It's the same point as behavior. I mean, you yourself said Skinner wasn't good because he was selecting the wrong things. That's because he was selecting behaviors. In fact, the things that are being selected for have more properties than just behavior. They have infinite classes of behavioral properties and the same thing with inputs.

GC: If you get to the basic phenomenon of what would be called control which is able to maintain your posture or your blood oxygen level or whatever in despite of disturbances. A servo-mechanism does capture it to some extent what's going on there at a simple variable level. So you have this phenomenon which is able to resist disturbances and you do that by having a reference level and a system where your response influences the perception and of course the perception influences the behavior, but, at the same time. So there's a phenomenon and this mechanism seems to capture some of the basics that is going on here. Okay, so there's no problem at that level, yet . . .

MB: Even at that level I would argue that it's a better characterization to
say that the servo-mechanism is "attempting" to achieve an internal state.

GC: An internal state? All it's really controlling is matching the input to the reference level and as you manipulate the reference level you can make an arm go up and do all kinds of things . . .

MB: But the result of that match is going to be some functionally efficacious internal state in the system and THAT'S the internal state it's trying to achieve.

GC: Not the reference level?

MB: Not the reference level. The only sense in which the reference level is the state it's trying to achieve--well, that's not even a state, it's a level--the only sense in which it's trying to achieve it is that under most normal working conditions, the matching of the reference level will achieve the internal state that it's after.

GC: Again, if I want to bring it back to a simpler level. If I have a cruise control on my car, I can manipulate the reference level and I can cruise at 65 or 55 or 45. You say that the system is not controlling the perceived speed coming back?. You're saying it's doing more than that?

MB: Of course it is. If that's all it did and that match didn't yield a further functional state, the servo-mechanism could not operate because it's that further functional state that either turns the system on or turns it off or switches to a different strategy or whatever it does. It's that further functional state that has all of the functional efficacies. The fact that there's a set point that under some conditions yields the further state is an additional point. You can have such a system with such a further state that doesn't involve a story of set points at all.

GC: That would resist disturbances?

MB: Sure. A TOTE mechanism does not require a comparator--that's just a simple fact. All you need is an internal functional state that serves as a switch and it either switches out of the system or it switches back to the operate in a TOTE. And that switch does not have to be an internal comparison switch. It's a switch that could be based on anything whatsoever.

GC: But the fact that you do see systems which are controlling what appear to be certain perceptions suggest that the switch is in fact operating.

MB: Some may be, right. And in some circumstances that will be adaptively desirable for them to be that way. But that's a limited class of circumstances.

GC: But in the phenomenon of control, you would recognize that organisms are controlling . . .

MB: Well, organisms are attempting to control their own internal states.

GC: And how do they know what their internal states are?

MB: Well, they don't know what they are but the internal states have functional efficacy, that's all.

GC: And what determines the functional efficacy?

MB: Just however it's organized. I mean, that has to do with the functional organization of the system so it's not so much that something determines it, it's rather that it's instantiated in some way or another in the nervous system. Now something will determine it in the sense of learning so then you need a model of how some part of a learning system can modify the function of those other parts of the learning system or some part of a learning system can modify its own function.

GC: So the notion of a higher level of controlled variables, something like
success as a professor or something like that, from a control-theory perspective, I would argue that you are going to perceive that and it may be some function of how you are dealing with your courses, how many students do you have, how much grant money is coming in, how many publications you're getting out, and if you perceive yourself as not matching that, you will adjust your behavior and vary it in someway in a hierarchical way to . . .

MB: At that sort of a level I would argue that thinking of that as a variable is very, very seriously distorted. It has a much richer structure than that. There are many different ways in which you can succeed as a professor. There are many different ways you can fail as a professor. Some are graded in the sense of ordered or partially ordered--some are not. Virtually none of them except salary has anything like a number line organization and so on.

GC: But you can see that it's appealing if you could extend this notion of hierarchy all the way up.

MB: Oh, I'm perfectly willing to do that. I would just argue that they don't all involve set points on a real number line. I would also argue that no version of that is ever going to get you to another knowing level. I would still further argue that a servo-mechanism hierarchy is not necessarily the most advantageous architecture for every possible tasks and that the brain doesn't necessarily use it for all possible tasks. I think there are reasons why it uses it for evolutionary common tasks like proprioceptive and kinesthetic control and so on, but I'm not at all persuaded that the brain necessarily uses servo-mechanism hierarchy architectures for higher level cognitive tasks.

GC: But are those tasks purposeful?

MB: Sure, but that doesn't mean that they're servo-mechanisms.

GC: So what is the other model then?

MB: Well, there's lots of those. For example, a lot of things in AI are written in what is called a blackboard architecture and the idea there is you have one big blackboard--sometimes you also have subordinate blackboards which give you different principles of hierarchy--and then you have a whole bunch of agents operating all at once and they're all showing their results on the blackboard and checking the blackboard to see if there are conditions on the blackboard that trigger their particular kind of activity. There's no servo-mechanism hierarchy there. So that's another sense in which the servo-mechanism architecture I would argue certainly does exist for some things and certainly can be used for some things but it's not a general architecture. It's not a general principle.

GC: But you would still consider that as being a purposeful system in some way? I'm trying to relate these blackboards with some task or some behavior or something like . . .

MB: But you can even construct a servo-mechanism out of the blackboard architecture. So if a higher level servo-mechanism throws a goal down to this blackboard and then all these little demons down here in parallel try to achieve the goal and under some condition this one's behavior would be more relevant than other conditions and these three over here will be relevant and so on, that's easy enough. But you don't have to think of it that way either. You can have the various demons in parallel throwing things onto the blackboard that in effect serves as goals for other demons.

GC: But they would be, in that case, higher in the hierarchy.

MB: Not necessarily because that can be a pure loop. There doesn't have to be any higher level. It can be a heterarchy not a hierarchy.

I think it's [a servomechanism hierarchy] an extremely powerful perspective. It's just not powerful enough. It's not a sufficiently general architecture. And when you try to apply it to things like I want to be successful as a professor, like I said, I just think it's highly distorted and I think it does as much harm as good at those sorts of levels because it's simply obscuring all of the structure there.

GC: It's trying to simplify it into a single variable somehow. And when I think about what some of these PDP circuits can do . . .

MB: But if you've got hierarchies of defeasibility relationships or hierarchies of critical principles, there's no way to construct a variable out of that. You cannot collapse into a variable the relationships of affirmation and infirmation and the defeasibility exceptions and all that kind of stuff. You can't do that with a real number line.

[Written comment by Mark Bickhard added when reviewing transcript}

MB: . . . it [the transcript] does not include the point that an automaton or Moore machine recognizer can serve the function of a functional test for a TOTE organization without there being any set point--a final state switches out--to Exit, and any other terminal state switches to Operate.
==

Date: Thu Jun 25, 1992 4:32 pm PST
Subject: Re: Good data, bad data

[Martin Taylor 920625 18:30]
(Bill Powers 920624.1500)

Sorry I can't reply properly to your postings in response to my flood of yesterday. I appreciate them, and will try to get to them over the weekend. Today, tomorrow, and Monday I have meetings all day.

But one point quickly...

>Just as Pat was about to start her PhD thesis using this experiment, BOTH
>SHE AND I UTTERLY CEASED TO HAVE ANY MOTION ILLUSIONS IN THIS EXPERIMENT.
>The offset of the illusions occurred about a week apart for us, mine
>disappearing first. As far as I can tell, nothing about the apparatus
>changed, although I beat my brains out trying to find some difference. Pat,
>incidentally, came out of this essentially cured of the debilitating motion
>disorder that had prompted her to do the study in the first place!

Wonderful! You were controlling the percept, weren't you? Have you ever had a motion aftereffect after driving a car? Actually, I'm not entirely sure that controlling is necessary for the aftereffect to disappear, since non-drivers learn not to see the aftereffect of forward motion. But I suspect that drivers lose it more quickly (but perhaps the control is of the zero motion encountered when you get out of the car and have to stand up; that's more like the experimental condition, and is the same for drivers and never-drivers). Check and see if you get an aftereffect of motion if you look out of the back window of the car for the length of a drive (with someone else driving).

Sorry to leave aside the interesting discussion on experimental method. I do hope to get to it, because I have an interesting blend of agreement and disagreement with you.

Martin

Date: Thu Jun 25, 1992 5:54 pm PST
Subject: Aftereffects; Bickhard blowhard

[From Bill Powers (920625.1900)]

Pat Alfano (920625) --

Didn't know you were monitoring the net, Pat. Hello, we're fine, nice to hear from you, hope the new house suits.

And CONGRATULATION ON THE DEGREE! WOW! YOU DONE IT ANYWAY!

You've filled me full of curiosity and embarrassed me, too. If the program didn't induce aftereffects in NEW subjects, then my whole fairy-tale goes down the drain -- including the "cure," I suppose. This is how people get hooked on organic vitamins and so on, isn't it? Well, you've given me a very good reason to redo that program -- obviously something DID change. But this is weird, because you will agree, won't you, that the display looks exactly the same as it did when there were aftereffects? I'm wondering now if shifting to a different monitor is what made the difference, although it's hard to imagine how, and I think we got aftereffects even on my new (VGA) monitor. I'm going to preserve the program that doesn't work and just develop the display again. If I get aftereffects again with the new program, then it will come down to comparing the programs in detail to see what's different. I'll let you know.

I seem to recall your telling me that you were having fewer difficulties with the motion stuff -- was that wishful thinking on my part? Of course that's unrelated to the disappearance of aftereffect, if new subjects failed to get it. If we're lucky we may find something nobody knew about that influences aftereffects. Or some dumb glitch in the program. Probably the latter.

Gary Cziko (920625) --

Fascinating talk with Bickhard. He gives the impression of knowing a very great deal about the higher orders of perception. That's a real skill. He should do well with grant proposals.

Just tell him there weren't any reactions from me that he (or his mother) would care to hear.

Pat Alfano has answered your question for me. Contact her through Dick Robertson.

Martin Taylor (920625.1830) --

>Wonderful! You were controlling the percept, weren't you?

That's right. See Pat Alfano's post for indications that this may be worth going on with (unless you've done it all before and have all the data). Watching a subject holding the dots "still," an observer could see that the dots kept moving, but more and more slowly. They moved, of course, in the original direction of movement, because when they stop, they seem to be drifting in the opposite direction and the subject "corrects" this apparent drift. An experiment in practical epistemology.

>Have you ever had a motion aftereffect after driving a car?

Only after long, long hours looking out the windshield; I think never unless I was generally pretty fatigued. When you stop the car, the road seems to be moving away. So does anything else you look at. But not after ordinary car trips (not even as a passenger).

I think, however, that your thesis might be upheld better in conjunction with walking or running. Driving a modern car isn't a lot different from being a passenger. And sometimes being the passenger isn't much different from being the driver.

I also tried using "endless octave" motion -- each dot would be placed randomly on the screen, move five or ten pixels on successive frames, and disappear. So the impression would be that the screen was crawling in some direction, but not actually going anywhere. My computer wasn't fast enough to do this with enough dots to get a strong effect.

Best Bill P.

Date: Thu Jun 25, 1992 8:10 pm PST
From: Dag Forssell / MCI ID: 474-2580
Subject: Leadership

[From Dag Forssell(920625)]

Hank Folson (920623)

An interesting perspective. Reminds me of "How to be a Christian in a Non-Christian world." A popular bible study group topic, I understand.

To promote PCT, our challenge is to show that knowledge of PCT can lead to greater satisfaction.

One aspect of this that I have noted is that (it seems to me) many people set impossible references for themselves: wanting to change history, their own upbringing (their subjective impression of it at that), or some aspect of their city or country that seems far removed from their capability to influence.

Continuing to consider the issue of following:

>Bruce Nevin (Mon 920422 15:02:32)
>
>>I can't conceive of a control system wanting to follow.
>
>Oh, come on, Dag! You can't mean that, can you?

Is there really such a thing as following? Or is this yet another instance of: "The human pie has already been sliced."

All we can do to each other as interacting control systems is to disturb one another (at the lowest levels.) The information we experience from a disturbance is interpreted and incorporated at the principle and system concept levels (if it is to have any leadership effect.) Each person controls him or her self given what s/he understands.

Alluding to my letter, let me suggest:

An understanding of cybernetic control (by which I meant PCT, of course) contains an explanation of the illusion of people following a leader.

Specifically, when Bruce Nevin leads in a dance, (you handsome hunk) an analysis of his follower's PURPOSES might suggest romantic systems concepts/principles causing the follower to control for exquisite, seductive "following."

When Saddam Hussein's followers' PURPOSES are analyzed, you might find some principle like survival of self and family. Just guessing.

It is ALL control!

Just like the meaning of "wife" is based on individual experience, so every word we use is subject to individual interpretation. LEADERSHIP is no exception. $175.- and the U.S. patent office says I get to define what I mean by Purposeful Leadership tm, and may try to explain it.

(This thread has been very helpful.)

A while back, I discovered that I had to re-evaluate my views on character issues in the light of PCT. I have accepted Rick Markens emphatic suggestion that: It is ALL control.

This places me in Hank's category 2.

Clearly, this process of having to re-evaluate our category 1 understanding afflicts all of us on this net. Bill included:

Bill Powers (920624.1200)

>I drift back and forth between paradigms, although as the years have
>gone by (more for me than most others), I've spent less time in the
>system concepts I was raised with and more in a new one. I think that
>when it comes to the crunch, I now opt for the new one almost all of the
>time.

It is amazing how many of the existing human pie slices need re-evaluation. (Bill suggested in the video Ed produced a few years ago that most people have a complete world view by age twelve, so anyone who did not learn PCT in kindergarten at the latest, will have a lot of re-evaluation to do.)

We are now looking at leadership. Sales is a very closely related subject, and then there is the framework of economics (where the sales take place,) which Bill has suggested for a thread.

The useful perspective seems to me to be to look at the follower's control, the buyer's control and of course the control processes involved in many individuals in the production, exchange and consumption of goods and services.

It occurs to me that "following" has both

1 a very technical meaning: direct employment by the follower of a signal as an adopted reference, which is what the dance partner does, which Bruce called me on, and

2 a "human pie slice" meaning: the attractive force the leader exerts on the follower to cause the follower to follow as a mindless victim, which I could not conceive of. "Charisma" comes to mind. Does it exist? No!

Reading and participating on the net sure provides an education. The recent discussion of statistics and the general lack of value of research based on a concept which is not at all descriptive of the phenomenon it is supposed to research seems in large measure a repeat of a year ago.

It is still fascinating to me, and a welcome refresher course. It seems obvious how difficult it is for a person to abandon cherished systems concepts, no matter how clear the arguments seem to others, who do not have the same personal investment in them.

Thomas Kuhn's "Scientific Revolution" is being played out here daily.

The show goes on. Remember: It is all control of perception.

Best, Dag.

Date: Fri Jun 26, 1992 9:26 am PST
Subject: Causation

[From Bill Powers (920626.0830)]

Mark Olson (920624) --

The apparent disagreement in definitions comes from my having thrown a knuckleball -- JOINT determinaton. If there are n influences on a variable, then the set of ALL N influences JOINTLY determines the behavior of the variable -- that is, we fully account for the behavior of the variable, knowing all of the influences on it. When we speak of influence, we mean only a partial accounting.

When I say that reference signals and disturbances JOINTLY determine actions, it's the plural of disturbance that makes this true. In effect, I'm saying that the reference signal plus all other influences on the controlled variable completely account for action -- i.e., determine it. Reference signal alone, or any one disturbance alone, can only INFLUENCE action.

But this way of speaking is misleading, because disturbances do not determine action in the sense of directly producing it. The disturbances are not brought to bear on the action, but on the controlled variable. The controlled variable does not directly affect the action, but just the opposite. What we're really saying is that the state of the controlled variable is fully accounted for by the action and the sum of all disturbances. GIVEN THE STATE OF THE CONTROLLED VARIABLE and the states of all disturbances influencing it, we can see that the action provides the missing influence that completes the explanation for the state of the controlled variable. So the state of the controlled variable is completely -- and jointly -- determined by the action plus the sum of all disturbances acting on it.

Of course we can't account for the state of the action without knowing the states of the controlled variable and the reference signal. So a correct picture of this situation can't come from a simple causal analysis. However we try to do it, there's always one variable unaccounted for, unless we consider the complete loop.

The concepts of causality and influence also omit something of critical importance about natural phenomena: the fact that the environment has PROPERTIES. The lever with a fulcrum in the middle has a property such that depressing one end by a given amount will, in general, raise the other end by a different amount. The ratio of movements is determined by the placement of the fulcrum. With the fulcrum in a given place, a movement that lowers one end by one inch will raise the other end by some different amount, perhaps 2 inches. But if we relocate the fulcrum, the same influence will raise the other end by a different amount -- half an inch. So the RELATIONSHIP between cause and effect is altered by moving the fulcrum, although if we simply move the fulcrum with no influence acting, there will be no behavioral effect.

You objected to my use of cause and effect in describing the lever, by saying

> Causal relationships must be temporal.

In fact, the relationship between movements of the ends of a lever IS temporal. The movement of the far end lags behind the movement of the near end by the time it takes a transverse wave to propagate from one end to the other, given the placement of the fulcrum. Just imagine a lever 20 feet long made of a 1/4-inch diameter steel rod. This delay is a consequence of other properties of the lever: its flexibility and its mass per unit length.

The concepts of causation and influence are left over from the time before modern science, before we thought in terms of a universe with properties. We still speak in such terms for the same reason we still speak of "looking at" objects, or "listening to" sounds, or "centrifugal force" -- an almost universal ignorance of physics. There are many other kinds of causation that we speak of because of ignorance of control theory.

I'm not saying that a translation from scientific terms into causal terms isn't useful. But don't rely on it for conveying a CORRECT picture of how things work. You can find an explanation of a Polaroid camera that will temporarily reduce the puzzlement of those who have never seen photography before, but until they get educated it can't be the right explanation (or at least one with fewer loopholes in it).

It was interesting to see that Birkhard, in conversation with Gary Cziko, didn't even understand how a cruise control works -- although he didn't hesitate to offer his explanation of, and rejection of, "servomechanism theory." He thought it worked by temporal cause and effect.

By the way, PV = nrt is an approximation to a physical law relating pressure, volume, and temperature. The actual physical law is far more complex than that, because if you add heat to an enclosed body of gas, the temperature will rise in a wave that diffuses through the gas; not all parts of the gas are at the same temperature. Only after the temperature has equalized throughout the volume will that equation hold (approximately) true. In the physical world, all processes take time to occur, and all laws of simple nature express only steady-state relationships.

Date: Fri Jun 26, 1992 9:50 am PST
Subject: Re: Leadership

[From: Bruce Nevin (Fri 920626 11:23:45)]

(Dag Forssell(920625)) --

>>Bruce Nevin (Mon 920422 15:02:32)
>>
>>>I can't conceive of control system wanting to follow.
>>
>>Oh, come on, Dag! You can't mean that, can you?
>
>Is there really such a thing as following? Or is this yet another
>instance of: "The human pie has already been sliced."

The original question was

  Can a control system want to follow another control system?

You are shifting now to ask another question:

  Can a control system follow another control system?

The "gather" program shows how one control system can follow another in terms of location. A control system follows another in this sense by setting as its goal a perception of proximity to the other.

Following surely cannot mean producing the identical behavioral outputs. We know this because of the variability of behavioral outputs with respect to the reference signal. (Or with respect to the outcome, more or less equivalent depending on success of control.)

Nor can it mean assuming the identical reference signals for identical (or equivalent) controlled perceptions. We know this because all the follower has to go by is the behavioral outputs of the leader, among other environmental variables. Plus memory and imagination, of course, which are the means for projecting, anthropomorphizing, and so on, which we necessarily do all the time.

There are two corresponding questions for the other member of the dyad:

Can a control system want to lead another control system?
Can a control system lead another control system?

From the existence of a large literature and a long history of "leadership" it seems clear that a control system can want to lead another.

It seems to me clear that A can lead B only to the extent and in the manner that B wants to follow A. This is why virtually all of traditional thinking about leadership boils down to "motivation"--getting others to want to follow you. (Ditto for pedagogy.)

Assume that B wants to follow A. The extent and manner depends on B's other goals. B can follow just in terms of proximity, as in the "gather" demo (puppy dog). This kind of following ranges from detailed minickry (mirroring) to very slight correlations such as B following A with his eyes.

Much of what we mean by "follow" is metaphorical, with this literal sense as a basis. We can easily identify the metaphor when we say B is "following A's argument" or "following A's line of thought."

The metaphor is not so obvious, perhaps, when we talk of B following A in the sense of coming to A for directions, going off and executing them, and coming back to A for more.

"Following directions" seems to mean to control one's perceptions so that they mimic ("follow") the perceptions that one imagines on hearing or reading the directions.

"Following A's argument" seems to involve imagining the argument for oneself, and finding that the imagined line of argument corresponds with what A has said and is saying.

>It occurs to me that "following" has both
>
>1 a very technical meaning: direct employment by the follower of a signal
> as an adopted reference, which is what the dance partner does, which
> Bruce called me on, and
>
>2 a "human pie slice" meaning: the attractive force the leader exerts on
> the follower to cause the follower to follow as a mindless victim,
> which I could not conceive of. "Charisma" comes to mind. Does it
> exist? No!

The first meaning concerns the question of how one control system can follow another. I have suggested some ideas about this above.

The second meaning concerns the question of why one control system would choose to follow another, the traditional question of how to "motivate" someone. I don't think this is a closed question at all. To paraphrase P.T. Barnum, some of the people want to follow all of the time, all of the people want to follow some of the time, but not all of the people want to follow all of the time.

If B is not confident and purposeful in a given situation, B may seek someone to follow until in a situation where B is more confident and purposeful. (Purposeful: has clear goals, is controlling for them without major conflict.) We can discuss why this is so.

If I am B in such a situation, I will follow one who appears confident and purposeful rather than one who appears unconfident and irresolute. We can discuss why this is so.

Some people are unconfident and irresolute and conflicted in much of their waking experience. I suspect that many such people came to be so because of childhood experience with adults who emphasized conformity with external authority and arbitrary standards, enforced in punitive ways.

It can happen that such a person feels confident and purposeful in an institutionalized social context with clearly assigned roles and relationships of relative authority, in accord with standards established for those institutions. Such people can become "leaders" within that framework. They know "the system." They become very anxious outside it, and resist contradiction to it. I think that outside the system they fear unexpected punishment; my experience is that outside the system (that is, in circumstances in which they can no longer interprete their perceptions as within the familiar institutional context) they become unconfident and irresolute. They often despise indecision and lack of confidence. (Such people, by the way, are unlikely to be drawn to HPCT at this stage in its history. And this parallels the familiar left/right ideological dichotomy.)

I suggest that charisma depends in part upon the appearance of confidence and purposefulness. As you have suggested, Dag, this connects with sales and marketing, where the pumped-up appearance often outstrips the basis of confidence and the real purposes are ulterior. But charisma can be genuine. When you're looking for the exit in a crowded waiting room, a person walking quickly in one direction with a suitcase has some charisma.

The ad hoc situational leadership and functional (not authoritarian) hierarchies of anarchism, as discussed in connection with Bill's story about V. the Boss, depend upon this, especially in cases where the participants lack detailed knowledge of another's capacities. "You seem to know what you're doing. How do you think we can make this go?"

Now:

Can a control system manipulate another control system?
Can a control system exploit another control system?

I believe these are some of the negative senses of "leadership" and "charisma" that you are resisting. Am I right? I think HPCT does not show that these do not exist. It only shows that they cannot work as intended. Social institutions can help people persist in being slow learners about this.

Bruce bn@bbn.com

Date: Fri Jun 26, 1992 2:14 pm PST
Subject: Teacher's reaction to PCT

from Ed Ford (920626:14:43)

To all -

Those attending the conference wishing letters inviting you or accepting your paper or whatever, let me know if I can help. I guess the current president of the CSG is also program chairman and I have sent several letters out already (Dick, your letter went out in today's mail).

For the past two weeks, I've been teaching elementary and high school teachers four hours a day, every day, for a professional education program at Ottawa U. here in Phoenix. I've never had a class so excited over control theory and its use in dealing with discipline, counseling problems (including parents), and stress. The first two days were somewhat quiet, then all of a sudden, the excitement level began to rise and rise. Their comments on the last day would make a control theorist's heart throb as they'd mention how this or that aspect of PCT made this or that part of their job or life easier to deal with. Obviously, their knowledge of PCT was limited, but I was very impressed with how much they understood and were translating into their own lives. I guess that's the kind of thing we're all trying to achieve.

Now to catch up on the volume of mail on the CSGnet.

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

Date: Fri Jun 26, 1992 3:05 pm PST
Subject: Re: Causation

[Martin Taylor 920626 17:45] -- meetings finished for today!
(Bill Powers 920626.0830 to Mark Olson)

Bill wrote such a lucid (as usual) comment, the I feel a little bit presumptuous trying to add to it. But I did think it would have been helped by another little diagram. And I'd like to use this diagram later in my follow-up on the statistics and experimental methods discussion.

```
                    To (from) higher levels
                   |        |
                   |        | Reference(s)
            --------\-------|-----------------
           |         \ comparator            |
           |          /          \           |
           |   perceptual        output      |  One of many ECSs at level n
           |    combiner        gain and     |
           |    function      distribution   |
           |   / / | | \ \     / / | | \ \   |
            --/-/--|-|--\-\----/-/--|-|--\-\--
            many control systems at many levels
   Person    ||||||||||||||    ||||||||||||||
             ^^^^^^^^^^^^^^    vvvvvvvvvvvvvv      Level 0 effects
   ========== many sensors ==== many effectors ==============
               \\\\\\\\\\         //////////      <--Physical observables
   World          \\\\\\           //////  <--Low-level interpretation
                     \\            // <--more abstract interpretation
                      ^            V <--A "solid" testable interpretation
                    controlled complex
                  environmental variable
                    //   \\   //  ||  \\
                   ^^     ^^ ^^   ^^   ^^
                 All sorts of disturbances
```

Note: The upgoing (sensor-based) flow and the downgoing flow (to the effectors) should be mingled, since reference signals tend to go down to low-level ECSs whose perceptual functions provide input upwards in a receprocal manner. But it is hard to draw that in ASCII.

The ECS controls one variable. To it, the variable is simple, mirrored in the single=valued result of its perceptual control function. But between that ECS and the variable it controls are many paths, both down through the effectors and back through the sensors. At each level down to the "world interface", there are other ECSs, each of which perceives a variable that is very simple from its own point of view. All these intermediate controlled variables and the ECSs that control them are part of the cause AND effect of the variable controlled by the ECS we are focussing on. All the percepts in the ECSs in the
control path mirror variables in the (excuse me) "cause-effect" circuit relating to the controlled variable in focus. And all are subject to disturbances that might or might not affect the focus variable.

In this diagram, it is hard to identify anything like causes, effects, influences, or properties. The nearest thing to a cause is the fact that changes in the reference at the top of the diagram cause mirroring changes in the environmental variable, if the properties and disturbances of the world permit it. But, if the reference stays stable, then the world disturbances can be said to "cause" restoring behaviour at the level marked with a single "V" marked "a solid testable interpretation" (which is both solid and testable, and more abstract than those above it).

The world disturbances "influence" behaviour that restores variables controlled by intermediate-level ECSs, since the reference levels that apply to those intermediate ECSs vary with the corrective outputs of the ECS in focus; hence there is no fully predictable effect of the disturbance on the intermediate- level behaviours. Consider the recent VOT discussion. If the talker was controlling for the perception that /p/ was being spoken, there is no need for VOT to be controlled except for being in the /p/ range, and it can therefore be partly controlled according to other references not determined in the experiment.

If we look very close to the world interface, the disturbances become "stimuli" and the controlled variables become "responses", but one can seldom talk about cause and effect at this level, firstly because the relevant references are unlikely to stay stable, and secondly because it is hard to observe either all the "stimuli" that combine in the perceptual function of any ECS or all the "responses" that might affect the variable that ECS controls.

The words "cause" "effect" and "influence" are very hard to give a precise meaning in a dynamic world that contains feedback. My description differs from Bill's in manner, but I think (hope) it carries the same import.

Martin

Date: Fri Jun 26, 1992 4:14 pm PST
Subject: Re: what's the difference?, psych research

[Martin Taylor 920626 19:40]
(Rick Marken 920624.1320)

>>We know that control can be exercised only to the extent that information is
>>available to the perceptual function of an ECS. And THAT is inherently a
>>statistical process,
>
>What is "THAT"? Information? The perceptual function? What is it about
>control that is "inherently statistical"?

THAT is the input to any ECS. Perception is largely a matter of extracting useful consistencies out of a very noisy sensory system sensing a highly variable world.

>
>>so we know that control IS a statistical process.
>
>When a non-statistical model accounts for 99.87% of the variance in
>the variables involved in control I think it's fair to say that the
>"inherently statistical" part of the process of control is not worth
>losing much sleep over.

True, in such an experiment. But remember that the amount of variance you account for depends on the ratio between the range over which the variable moves and the size of the unaccounted variation. Even in a tracking study, if the target moved only over a range of 1 mm on a screen viewed at a normal distance, I doubt you would find 99% correlations anywhere in your analysis. In that kind of study, you can make the range of variation much larger than the statistical variability, and good, more power to you. But I doubt you can do it so readily at higher levels or under more noisy conditions.

In the psychoacoustic experiments you so often exclaim against, the whole problem is the determination of the perceptual variability. There can be no PCT-based study in which control will do better than a perfect ECS whose perceptual function is a mathematically ideal observer. Humans, well trained, can come within 3 or 4 dB of that, under a wide variety of conditions. Perceptual statistical variability has to be a limiting condition for control, and hence control is inherently a statistical process.

Choose your experiment so that perceptual noise is swamped by big disturbances in the controlled variable, and if none of the other factors I mentioned in [920624 14:10] is important, then you may get your high correlations.

>Go for the QUALITY data.

Yes, the best that suits the problem at hand. And I grant that PCT experiments
are likely to do better than non-PCT experiments, for good reason.

More later, on the other "statistical" postings.

Martin

Date: Fri Jun 26, 1992 6:44 pm PST
Subject: Taylor's diagram!!!!

[From Bill Powers (920626.1900)]

Martin Taylor (920626.1745) --

What a stunning exposition on HPCT. Looking at your diagram, I immediately saw what's wrong with my remarks to Mark Olson on "properties". What constitutes a property of the external world depends on what level of control you're talking about. The observer, trying to find a controlled quantity, views the fine-grained environment through perceptual systems that are hoped to be like those in the organism under study. The environment itself, therefore, is represented in a way suitably abstracted for the way behavior is being investigated. And naturally, the properties to be found in that environment are different from what they would be if one were looking at a different level of control.

Higher-order properties of the world are almost certain to be defined in terms of descriptions of plant, animal, and human participation in that world, because the world as described by physics doesn't "have" higher- order properties like categories and so on up. And anyway, physics describes the world without control by plants and animals and human beings -- it is the science of Mars or Venus, but not of Earth.

Remember that silly bit I wrote about "truthsaying?" I think you got into that mode. I feel that something has been illuminated for me.

I think that what you've put together should become part of our basic introduction to control theory. I'll go farther: I think you ought to try to develop this picture into a serious article, aimed at making sense of HPCT for the scientific community. I think that this presentation will have a powerful effect. So far -- on third reading -- it's flawless.

I've been trying for a long time to get all the factors in your diagram put together into a single neat package. You've done it. I have the feeling that when I wake up tomorrow and look at it again, it's going to look just as good. It's as though you got into my messy head and said "No, no, this goes HERE and that goes THERE, and you can leave out that other junk. Isn't this what you mean?" It sure is. I feel like celebrating. Can anyone tell?

Martin, to save a little time, a comment on statistical perception. You said privately that if one swamps a controlled variable with a large enough disturbance, then experiments will show high correlations. I think we're on the verge of an agreement here. My thesis is simple: the normal range of perceptual signals in normal control behavior is such that the signal-to-noise ratio is very high; I would guess between 10:1 and 100:1. When perceptions begin to approach their lower limits of magnitude, then the impulse nature of the neural signal becomes important, just as it does in electronics. Is this going to solve our problem?

Rick: Roger.

Best to all, Bill

Date: Sat Jun 27, 1992 9:08 am PST
Subject: Misc

[From Rick Marken (920627.1000)]

Unfortunatly I deleted Joel Judd's recent post on psych research (by accident). But it was excellent. It's a good thing you already have your PhD; it's not nice to tell your committee members that their fundamental assumptions are a crock.

Speaking of fundemental assumptions that are a crock, how about Gary's interview with Mark Bickhard -- Luce professor, no less (is that Henry or R. Duncan?) of more things than I can think about at the same time. My hat's off to Gary for sitting though the whole thing -- let alone asking pertinent questions at the same time. I used to think that you had to go to the literature on "dynamic attractor" models of behavior to get that level of high-falutin' incoherence. I don't think Dr. Bickhard would enjoy CSGList any better than did Dr. Beer. I think I'll go have a Dr. Pepper.

I'll be back.

Rick

Date: Sat Jun 27, 1992 9:23 am PST
From: Dag Forssell / MCI ID: 474-2580
Subject: Leadership

[From Dag Forssell (920627-1)]

Bruce Nevin (Fri 920626 11:23:45)

>The original question was
> Can a control system want to follow another control system?
>You are shifting now to ask another question:
> Can a control system follow another control system?

When I said that

>>>>I can't conceive of control system wanting to follow.

I did not mean to be so literal. I meant that it is not the nature of a control system to "follow", whatever that is.

When I said that

>>Is there really such a thing as following? Or is this yet another
>>instance of: "The human pie has already been sliced."

again, I did not mean to be quite so literal as your paraphrasing. I appreciate your post. Your restating my points, paraphrasing rather, is a very good thing. It shows me how my careless wording can be (mis-) interpreted. You are doing a good job of sorting out technical alternatives and aspects of "following."

>Now:
>
> Can a control system manipulate another control system?
> Can a control system exploit another control system?
>
>I believe these are some of the negative senses of "leadership" and
>"charisma" that you are resisting. Am I right? I think HPCT does not
>show that these do not exist. It only shows that they cannot work as
>intended. Social institutions can help people persist in being slow
>learners about this.

You are hitting the nail on the head here. The answer to these questions is an obvious yes. I agree that these exist. A question that then arises is: What do we mean by "manipulate" and "exploit". You have just participated in a discussion on issues of stereotyping; projecting prejudices on others based only on your own subjective experience. I am resisting what I perceive to be extremely common stereotype interpretations of leadership and sales, where I sense an interpretation that leadership and sales are indeed "manipulation" and "exploitation". This I read into Bills original refusal to lead and some comments about sales at past conferences.

In turn, this leads to an aversion to consider these major applications of HPCT. Still my perceptions, of course.

If you substitute "manipulate" with "inform", "guide" "enlighten" "teach" and "exploit" with "mutual benefit," the substance of the interaction does not change from a HPCT point of view, but the emotional, stereotype flavor changes dramatically.

We are still talking about leadership and sales and mutual economic advantage.

Certainly the members of this net want to sell HPCT to the world. Is this "manipulation" and "exploitation?" I would not label it that. But mention leadership and sales. What comes to mind? Some brutal, selfish "leader" on the one hand and pusher of overpriced junk nobody wants or needs on the other.

These terms are among the unexamined "human pie slices" - systems concepts from pre-HPCT days - that can benefit from some HPCT light. By looking closely at this, perhaps a way to sell HPCT can be found, vastly superior to the frustrating sales efforts into the psychological journals that are discussed here, but not labeled as such. (These journals are a minuscule market compared to the rest of the world and the one market where we know that PCT is not welcome).

The way there is to forget about "manipulation" and "exploitation" and instead examine the best interest of and control processes in the other autonomous control system, whether we call it follower or buyer. This done with full visibility to said follower and buyer, of course. There is no need to hide the interest and control processes of the leader or salesperson either. The exchange of goods or services should benefit both parties. Otherwise we have reverse manipulation and exploitation.

Leadership and sales both can be honorable. HPCT can show how.

Again, I find it useful to be challenged to think this through again myself.

Best to all. Dag.

Date: Sun Jun 28, 1992 12:29 pm PST
Subject: Rock & Roll [From Rick Marken (920628)]

Well, now I know why LA is the rock and roll capital of America. Two earthquakes this morning. The first was a doosy. I estimated it at 7.2. The last I heard the instrument based measure was 7.4. The second was a bit smaller (by my perception) - I estimated 7.0 I heard that that was the current estimate.

Talk about perceptions that you can't control. I don't remember feeling a nice, sustained quake like that since I was a kid (and, indeed, they said that this was the biggest earthquake, in epicentral magnitude, to hit the LA area in 40 years. And I remember THAT one. I was six and my dad took me for a tour of the house during it; he's a phenomena freak too).

I have been able to call earthquake magnitudes to within 0.2 Richter points for the last four earthquakes that I have experienced. This is quite a feat (if I do say so myself) since my distance from the epicenter is, of course, always quite different (I was as close as 10 miles to one; I was a comfortable 120 mi from this one). I wonder if anyone has done any earthquake psychophysics?

Have a nice, stable sunday. Rick
Date: Sun Jun 28, 1992 12:53 pm PST
Subject: AI, Robots, & Rodney Brooks

Hello-

First, an introduction: I'm a systems programmer and have been in computer science study/browse mode for a couple of years. As motivation, I've been writing a 3D real-time (quasi virtual-reality) window system.

I've recently discovered Rodney Brooks' work at the MIT AI Robotics' Lab. If you folks aren't familiar with this, you should be. He is leading a reaction to AI's traditional robotics goals: he wants to build non-cognitive robots that nevertheless do interesting things in their environments, and he uses insects as models instead of humans.

His approach is to study insects, snails, and other very simple creatures that survive in complex environments, and steal their design architecture. A traditional robot architecture looks like this:

```
          planning layer
      / \              |
       |              \ /
    sensors       motor control
```

where "planning layer" means huge amounts of raw, seething LISP on a separate computer with large cables between it and the robot.

Brooks' approach looks like this:

```
         -> behavior 3               ->|
         -> behavior 2b        ->|     |
         -> behavior 2a  ->|     |     |
 sensors -> behavior 1  -> X  -> X  -> X  -> motor control
```

In this architecture, a 'behavior' is a simple control system box. An X is an override box. Each behavior is active or inactive. If it decides it's precondition occurs, it countermands the orders of lower behaviors. In the case of a robot which wanders around picking up empty soda cans and takes them back to a "home base", behavior 1 attempts to move forward, behavior 2a sees anything ahead and attempts to turn left, 2b sees anything ahead and to the left and tries to turn right, and 3 sees something ahead that looks like a can and tries to stop. There is no hierarchy of information flow, because that is too slow. Each layer interprets the same raw data according to its own mission.

Behaviors are implemented with very little computer power. Microcontrollers with EPROMS are the preferred method; they've got radio-controlled cars with self-contained CPU boards replacing the radio control units. These things are doing grouping behavior (i.e. birds in Vs, fish in schools).

I'm personally interested in this work because I want a method of implementing interesting autonomous non-boring robots in my VR window system without hooking up a Cray.

Anyway, Brooks' work attempts to start towards consciousness/ intelligence by building gangs of control systems together with clever design. He's going at consciousness from the bottom up with the same architecture that you folks are using from the top down.

Here is a bibliography of technical reports available from the MIT AI Lab. "Intelligence Without Reason" is available via FTP as ftp.ai.mit.edu:pub/doc/brooks-ijcai91.ps.Z pub/publications on that same machine has the full list of technical reports that this list is culled from. Also, the John Connely book detailing the above soda-can-picking robot is also in print from the MIT press in hardback for $30 instead of $13.

Maybe you can put this up for FTP instead of posting it, I don't know what your policies are.

% Updated January 1992
% Please check the pub/publications directory for other information
% listed on-line. The file is pubs-online

% PLEASE NOTE: It is now necessary to chrage for shipping. At present
% we will include the shipping charges with your order and we ask that
% you pay them at your earliest convenience. Thank you.

% TO ORDER, specify publications number and author and enclose
% a check payable to the M.I.T. Artificial Intelligence Laboratory for
% the correct amount of U.S. funds.

% PREPAYMENT IS REQUIRED. Please note on order if check is sent
% separately.

% Send orders with payment to:

% Publications, Room NE43-818
% M.I.T. Artificial Intelligence Laboratory
% 545 Technology Square
% Cambridge, MA 02139 USA

% For additional information:

% Phone number: (617) 253-6773}
% Net address: Publications@ai.mit.edu}

% Another source for these publications is:

% NTIS. Reports assigned an ``AD'' number (such as AD-A123456) are
% available from the National Technical Information Service, 5285 Port
% Royal Road, Springfield, Virginia 22161. NTIS price information can
% be obtained by calling (703) 487-4650.

% If you would like to receive future updates by e-mail,

% please send your net address to publications@wheaties.ai.mit.edu

% This is the master list of current AI memos and technical reports.
% The format is for the LISP program that generates the bibliography.
% Memos are listed first, then technical reports. To jump directly to
% the TRs, search for ":tr 474".

:aim 842
:author Tom\'as Lozano-P\'erez and Rodney A. Brooks
:asort Lozano-P\'erez, T.; Brooks, R.A.
:title An Approach To Automatic Robot Programming
:date April 1985
:pages 35
:cost $4.00
:adnum AD-A161120
:keywords robotics, task planning, robot programming
:abstract
In this paper we propose an architecture for a new task level system, which we call TWAIN. Task-level programming attempts to simplify the robot programming process by requiring that the user specify only goals for the physical relationships among objects, rather than the motions needed to achieve those goals. A task-level specification is meant to be completely robot independent; no positions or paths that depend on the robot geometry or kinematics are specified by the user. We have two goals for this paper. The first is to present a more unified treatment of some individual pieces of research in task planning, whose relationship has not previously been described. The second is to provide a new framework for further research in task-planning. This is a slightly modified version of a paper that appeared in {\it Proceedings of Solid Modeling by Computers: From Theory to Applications}, Research Laboratories Symposium Series, sponsored by General Motors, Warren, MI, September, 1983.

:aim 864
:author Rodney A. Brooks
:asort Brooks, R.A.
:title A Robust Layered Control System For a Mobile Robot
:date September 1985
:pages 25
:cost $3.00
:adnum AD-A160833
:keywords mobile robot, robot control
:abstract
We describe a new architecture for controlling mobile robots. Layers of control system are built to let the robot operate at increasing levels of competence. Layers are made up of asynchronous modules which communicate over low bandwidth channels. Each module is an instance of a fairly simple computational machine. Higher level layers can subsume the roles of lower levels by suppressing their outputs. However, lower levels continue to function as higher levels are added. The result is a robust and flexible robot control system. The system is intended to control a robot that wanders the office areas of our laboratory, building maps of its surroundings. In this paper we demonstrate the system controlling a detailed simulation of the robot.

:aim 899
:author Rodney Brooks
:asort Brooks, R.A.
:title Achieving Artificial Intelligence through Building Robots
:date May 1986
:pages 12
:cost $3.00
:adnum AD-A174364
:keywords artificial intelligence, robotics
:abstract
We argue that generally accepted methodologies of artificial intelligence research are limited in the proportion of human level intelligence they can be expected to emulate. We argue that the currently accepted decompositions and static representations used in such research are wrong. We argue for a shift to a process based model, with a decomposition based on task achieving behaviors as the organizational principle. In particular we advocate building robotic insects.

:aim 984
:author Rodney A. Brooks, Anita M. Flynn, and Thomas Marill
:asort Brooks, R.A.; Flynn, A.M.; Marill, T.
:title Self Calibration of Motion and Stereo Vision for Mobile Robot Navigation
:date August 1987
:pages 25
:cost $3.00
:adnum AD-A185602
:contract N00014-86-K-0685, IBM, SDF, N00014-K-0124
:keywords mobile robot, self calibration, stereo vision, motion vision
:abstract
We report on experiments with a mobile robot using one vision process (forward motion vision) to calibrate another (stereo vision) without resorting to any external units of measurement. Both are calibrated to a velocity dependent coordinate system which is natural to the task of obstacle avoidance. The foundations of these algorithms, in a world of perfect measurement, are quite elementary. The contribution of this work is to make them noise tolerant while remaining simple computationally. Both the algorithms and the calibration procedure are easy to implement and have shallow computational depth, making them (1) run at reasonable speed on moderate uni-processors, (2) appear practical to run continuously, maintaining an up-to-the-second calibration on a mobile robot, and (3) appear to be good candidates for massively parallel implementations.

:aim 1016
:author Rodney A. Brooks, Jonathan H. Connell, and Peter Ning
:asort Brooks, R.A.; Connell, J.H.; Ning, P.
:title {HERBERT: A Second Generation Mobile Robot}
:date January 1988
:pages 11
:cost $3.00
:adnum AD-A193632
:contract N00014-86-K-0685, SDF, N00014-85-K-0124
:keywords mobile robot, parallel processor, laser scanner

:abstract
In mobile robot research we believe the structure of the platform, its capabilities, the choice of sensors, their capabilities, and the choice of processors, both onboard and offboard, greatly constrain the direction of research activity centered on the platform. We examine the design and tradeoffs in a low cost mobile platform we have built while paying careful attention to issues of sensing, manipulation, onboard processing, and debuggability of the total system. The robot, named Herbert, is a completely autonomous mobile robot with an onboard parallel processor and special hardware support for the subsumption architecture [Brooks (1986)], an onboard manipulator, and a laser range scanner. All processors are simple low speed 8-bit micro-processors. The robot is capable of real time three-dimensional vision, while simultaneously carrying out manipulator and navigation tasks.

:aim 1091
:author Rodney A. Brooks
:asort Brooks, R.A.
:title A Robot That Walks; Emergent Behaviors from a Carefully Evolved Network
:date February 1989
:pages 12
:cost $3.00
:adnum AD-A207958
:contract N00014-86-K-0685, N00014-85-K-0124
:reference Also in {\it Neural Computation}, vol. 1, no. 2, 1989.
:keywords subsumption architecture, walking robot, emergent behavior, distributed control
:abstract

This paper suggests a possible mechanism for robot evolution by describing a carefully designed series of networks, each one being a strict augmentation of the previous one, which controls a six- legged walking machine capable of walking over rough terrain and following a person passively sensed in the infrared spectrum. As the completely decentralized networks are augmented, the robot's performance and behavior repertoire demonstrably improve.

:aim 1120
:author Anita M. Flynn, Rodney A. Brooks, William M. Wells III, and David S. Barrett
:title SQUIRT: The Prototypical Mobile Robot for Autonomous Graduate Students
:date July 1989
:pages 31
:cost $4.00
:adnum AD-A212337
:contract N00014-86-K-0685, N00014-85-K-0124
:keywords miniature robot, autonomous robot, subsumption architecture
:abstract
This paper describes an exercise in building a complete robot aimed at being as small as possible but using off-the-shelf components exclusively. The result is an autonomous mobile robot slightly larger than one cubic inch which incorporates sensing, actuation, onboard computation, and onboard power supplies. Nicknamed Squirt, this robot acts as a ``bug,'' hiding in dark corners and venturing out in the direction of last heard noises, only moving after the noises are long gone.

:aim 1126
:author Anita M. Flynn, Rodney A. Brooks, and Lee S. Tavrow
:asort Flynn, A.M.; Brooks, R.A.; Taurow, L.S.
:title Twilight Zones and Cornerstones: A Gnat Robot Double Feature
:date July 1989
:pages 44
:cost $4.00
:adnum AD-A220026
:contract N00014-86-K-0685, N00014-85-K-0124
:keywords gnat robot, micro robot, piezoelectric motor, IR/Optical camera, recursive gnat robot assembly line, disposable robots
:abstract
We want to build tiny gnat-sized robots, a millimeter or two in diameter. They will be cheap, disposable, totally self-contained autonomous agents able to do useful things in the world. This paper consists of two parts. The first describes why we want to build them. The second is a technical outline of how to go about it. Gnat robots are going to change the world.

:aim 1148
:author Anita M. Flynn and Rodney A. Brooks
:asort Flynn, A.; Brooks, R.
:title Battling Reality
:date October 1989
:pages 21
:cost $3.00
:adnum AD-A220016
:contract N00014-85-K-0685, N00014-85-K-0124
:keywords subsumption architecture, mobile robots, gnat robots, robotics, sensors, navigation
:abstract
In the four years that the MIT Mobile Robot Project has been in existence, we have built ten robots that focus research in various areas concerned with building intelligent systems. Towards this end, we have embarked on trying to build useful autonomous creatures that live and work in the real world. Many of the preconceived notions entertained before we started building our robots turned out to be misguided. This paper describes the changing paths our research has taken due to the lessons learned from the practical realities of building robots.

:tr 1151
:author Jonathan Connell
:asort Connell, J.H.
:title A Colony Architecture for an Artificial Creature
:date August 1989
:pages 131
:cost $8.00
:adnum AD-A216802
:contract N00014-85-K-0124, N00014-86-K-0685
:keywords subsumption, robotics, mobile robot, multi-agent, autonomous, collection
:abstract
This report describes a working autonomous mobile robot whose only goal is to collect and return empty soda cans. It operates in an unmodified office environment occupied by moving people. The robot is controlled by a collection of over 40 independent ``behaviors'' distributed over a loosely coupled

network of 24 processors. Together this ensemble helps the robot locate cans with its laser rangefinder, collect them with its on-board manipulator, and bring them home using a compass and an array of proximity sensors. We discuss the advantages of using such a multi-agent control system and show how to decompose the required tasks into component activities. We also examine the benefits and limitations of spatially local, stateless, and independent computation by the agents.

:aim 1182
:author Rodney A. Brooks and Anita M. Flynn
:asort Brooks, R.; Flynn, A.
:title Fast, Cheap and Out of Control
:date December 1989
:pages 14
:cost $3.00
:contract N00014-86-K-0685, N00014-85-K-0124
:keywords autonomous rovers, legged robots, gnat robots, subsumption architecture, space exploration
:abstract
Spur-of-the-moment planetary exploration missions are within our reach. Complex systems and complex missions usually take years of planning and force launches to become incredibly expensive. We argue here for cheap, fast missions using large numbers of mass produced simple autonomous robots that are small by today's standards, perhaps 1 to 2kg. We suggest that within a few years it will be possible, at modest cost, to invade a planet with millions of tiny robots.

:aim 1227
:author Rodney Brooks
:asort Brooks, R.
:title The Behavior Language: User's Guide
:date April 1990
:pages 35
:cost $4.00
:adnum AD-A225808
:contract N00014-86-K-0685, N00014-85-K-0124
:keywords subsumption, behavior language
:abstract
The Behavior Language is a rule-based real-time parallel robot programming language originally based on ideas from [Brooks 86], [Connell 89], and [Maes 89]. It compiles into a modified and extended version of the subsumption architecture [Brooks 86] and thus has backends for a number of processors including the Motorola 68000 and 68HC11, the Hitachi 6301, and Common Lisp. Behaviors are groups of rules which are activatable by a number of different schemes. There are no shared data structures across behaviors, but instead all communication is by explicit message passing. All rules are assumed to run in parallel and asynchronously. It includes the earlier notions of inhibition and suppression, along with a number of mechanisms for spreading of activation.

:aim 1230
:edited by Anita Flynn
:asort Flynn, A.
:title Olympic Robot Building Manual
:date December 1988

:pages 118
:cost $6.00
:keywords

:aim 1293
:author Rodney A. Brooks
:asort Brooks, Rodney A.
:title Intelligence Without Reason
:date April 1991
:pages 29
:cost $3.00
:contract
:adnum AD-A241158
:keywords artificial intelligence, situatedness, embodiment
:abstract
Computers and Thought are the two categories that together define Artificial Intelligence as a discipline. It is generally accepted that work in Artificial Intelligence over the last thirty years has had a strong influence on aspects of computer architectures. In this paper we also make the converse claim; that the state of computer architecture has been a strong influence on our models of thought. The Von Neumann model of computation has lead Artificial Intelligence in particular directions. Intelligence in biological systems is completely different. Recent work in behavior-based Artificial Intelligence has produced new models of intelligence that are much closer in spirit to biological systems. The non-Von Neumann computational models they use share many characteristics with biological computation.

Date: Sun Jun 28, 1992 4:03 pm PST
Subject: Beer in the Brooks

[From Rick Marken (920628.1600)]

Martin Taylor's diagram:

This is, indeed, a nice way of showing that the environment part of a hierarchical control system is a hierarchical structure to the extent that the systems observing that environment are also hierarchical control systems. I tried to make the same point, verbally, in my "Hierarchical behavior of perception" paper, that everyone is rushing to avoid publishing.

The diagram is particularly relevant, I think, to the observation of human behavior (since the behaviors that we see controlled by people are likely to correspond to perceptions we can have, even if we can't control them -- though this is not necessarily the case; it is possible to imagine watching a human behave and have no idea what s/he is controlling. I think this happens when we watch certain skilled behaviors -- like that of a surgeon. It's often hard to see what aspect of the surgeon's environment is controlled. There can also be intentional obfuscation of a controlled variable -- as in magic.)

With organisms other than humans, many of the variables that are controlled by the organism are not normally part of the human perceptual hierarchy -- like the high frequency chirps of the bat. Nevertheless, these variables can be detected by artificial sensors -- but all the human really perceives in this case is the "meter reading" as the controlled variable.

Technically Sweet <thinman@NETCOM.COM> says:

>Hello-

Welcome to CSG-list, Mr. Sweet.

>I've recently discovered Rodney Brooks' work at the MIT AI Robotics'
>Lab. If you folks aren't familiar with this, you should be.

We are familiar with him.

>He is leading a reaction to AI's traditional robotics goals:

We (some people on CSGNet) just had a long interaction about the work of a fellow named Randall Beer whose appoach seems similar to Brooks' (he makes insects, uses "independent behavior producing agents", etc -- I think the term "subsumption architecture" is TM Brooks -- but it looks like Beer does the same thing). Beer and Brooks seem completely uninterested in PCT (our model) apparently because they are modelling a different phenomenon than we are. They are modelling "behavior" which to them means certain observable "outputs" generated by the organism -- like "finding soup cans" and "collecting them". We are interested in modeling "control" -- the ability or organisms to keep aspects of their own perceptual experience in internally specified states (the specifications may be fixed or variable). "Control" and "output generation" LOOK the same (on casual inspection) but they are not. We have spent PAGES trying to communicate the difference -- so don't expect it here in a sentence. If you're interested, stay tuned to CSGNet.

The architectures we use to model control are completely different than those used by Brooks to model behavioral output. The focus of our models is the design of the perceptual function and the design of methods for specifying the intended level of the outputs of these functions.

>I'm personally interested in this work because I want a
>method of implementing interesting autonomous non-boring
>robots in my VR window system without hooking up a Cray.

Brooks's is probably as good an approach as any if that is your goal (incidentally, I passed a toy store yesterday and saw some nice little "robot" puppies barking away in a cage -- poking their noses through the bars as though they were trying to lick the hands of passersby. I could barely resist the urge to pet them. These dogs were not controlling anything -- but they were exhibiting some pretty interesting behavior. I'm sure the behavior exhibited by Brooks' insects could be even more "non-boring" -- but it is mostly NOT control -- though some may be because at the lowest levels of these architectures there is a closed loop).

It is not easy to understand the difference between behavior as an interesting display of "outputs" and behavior as "control". That is one of the main things CSGNet is about. It is a VERY difficult concept to get -- and it is virtually impossible to get if you don't believe there is such a distinction in the first place. Beer seemed quite uninterested in PCT; I believe Brooks is similarly inclined. They seem to be happy (and famous too, at least Brooks) and quite disinclined to discuss the issue. More power to them.

> He's going at consciousness from the
>bottom up with the same architecture that you folks are
>using from the top down.

Actually, we're going at "purpose" (control) -- usually from the inside out but we're trying to get at from the outside in too. We're not really very close to putting "consciousness" into the model yet.

Best regards Rick

Date: Sun Jun 28, 1992 9:38 pm PST
Subject: PCT REVOLUTION

From Hank Folson (920629)

Nothing like a cooouple of 'quakes to remind me to get on the Mac & respond to Bill's comments. I agree with everything you say, Bill, especially where you clarified what I was trying to say:

>[From Bill Powers (920624.1200)]
>>Hank Folson (920623) --

>>Wouldn't the controlling actions be different?
>Not so much the actions as the goals.

Your choice of words is clearer than mine: "Wouldn't the _goals_ be different?" is what I meant to say.

>>1A...>>techniques that are compatible with PCT?

>All methods of control are compatible with PCT...

Again, I was not clear. Is this better: 1A.(A small but important sub-category) How will a person (living control system) control when s/he does not know that s/he is a living control system, and does not know that everyone else is an independent control system that cannot be controlled without force, but has through their life experience developed a Systems Concept, _goals, and controlling actions that a person who understands and applies PCT would?_

>>Would you gain any new insights by using these categories

>Probably. But don't be shy: tell us what new insights you think we might
>get.

It's not shyness, Bill. I'm controlling to learn, not to teach (As I think about it, one can do both.) so I did not offer my interpretation of reality. Here is my take on Dag Forssell's (920621-1) jury verdict article about the police charged with killing a drug dealer, with strong evidence against the police:

>Richard, a 38-year-old engine mechanic, said (during deliberations) that
>Mercado was "only a drug dealer, anyway."
>Herbert, a 59-year-old airline mechanic, believed that "criminals give
>their civil rights away when they elect to lead a life of crime."

These statements reflect the speakers' System Concept that the police are good people serving the community and drug dealers are bad people. The dead man fit their category of bad people, the police fit their category of good people. These categories are very strong, and long held, so these jurors would not perceive the evidence as proof that in this instance, the roles may have been reversed. Control systems resist change.

>Rubina, a 53-year-old saleswoman, didn't believe several prosecution
>witnesses from the neighborhood because "these are the people we're
>paying the policemen to protect us from."

Rubina's system concept is that all the people in this part of town are all in the bad people category.

>Bernie, a 48-year-old butcher, thought the police were guilty, but he
>changed his vote because "I didn't want to be the one that was sitting
>out there with them pointing at [me]."

Bernie is controlling for acceptance and belonging. Fitting in is more important to him than making an independent evaluation of the evidence.

>Most telling, perhaps, was one juror's observation that the officers had
>to be found guilty "beyond an absolute doubt." This juror had
>single-handedly changed the standard of doubt in a criminal case. I
>suspect he did so because he felt more sympathy for police fighting the
>drug war than for a drug dealer with a violent past.

Here the author's perception appears accurate, from a PCT point of view.

As has been said many times on this 'net, it's all perception. Had these jurors been raised in an environment in which truth and justice were very important aspects of their System Concepts, their perceptions of the evidence would have been different. It is possible that years ago, before drugs and crime became such a big problem (real or perceived) in their society, these same people might have reached a different verdict. But if the media, War On Drugs, their awareness of friends and relatives suffering because of drugs, etc., affected them strongly enough and often enough, their System Concepts may have evolved to where they were at the time of the trial.

The jurors are all Category 1 people, unaware we are all control systems who will control our perceptions. They have no understanding of levels, especially those involved here, which include Category, Principle and System Concept. So they have no awareness that they are choosing the way they perceive the police and the drug dealer/victim. If they did, they might be able to separate their own System Concept from a more generalized and higher concept of how their society should handle this situation, and choose which way to go. It is interesting how open they all were in describing what they were controlling for.

The next aspect to consider is the Systems Concept of the author. He is controlling for something, even though he, too, is Category 1. We do not know what quotes he chose not to include in the article. Reading his complete notes could tell a lot about his view of the world. Why did he choose these particular quotes? The whole article is presented through a filter. We are only seeing the results of what the author is controlling for, if the article is our only source of information.

As I read the rest of the article, the author is in internal conflict. On one hand, he knows that by his System Concept, these police were guilty as Hell. On the other, he has an uneasy awareness and understanding of why the jurors did not give a guilty verdict. He knows pretty much what is going on, he just doesn't know why. If he understood PCT, he could understand where he was coming from and where the jury was coming from. And if he did, he might have written an article that showed less internal conflict, explained confidently what was going on, etc.

Not only are you getting this information filtered by the author, you are getting it filtered again by Hank Folson (a Category 2 person, but that's my perception). What am I controlling for? Whatever it is, it has affected the way I have presented this.

What are you controlling for when you read this? You are doing your own personal filtering of reality as well. Who knows what is getting through to you about this case and about PCT?

>RE my quote: don't overlook the part that says "... if we want to find >the significance of the first new concept of human nature since >Descartes." I'm trying to say what we need to do IF that's the goal. For >a lot of people, it still isn't.

Count me in, Bill, if you're controlling for what I think you're controlling for.

Trustingly, 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: Mon Jun 29, 1992 8:25 am PST
Subject: Teaching system concepts; Brooks & PCT

[From Bill Powers (920629.0930)]

Hank Folson (920629) --

Your post on various ways of handling knowledge/ignorance about control is most thoughtful. It might point toward some educational reforms (as long as everyone's talking about them again, for the election year).

>The jurors are all Category 1 people, unaware we are all control
>systems who will control our perceptions. They have no understanding of
>levels, especially those involved here, which include Category,

>Principle and System Concept. So they have no awareness that they are
>choosing the way they perceive the police and the drug dealer/victim.
>If they did, they might be able to separate their own System Concept
>from a more generalized and higher concept of how their society should
>handle this situation, and choose which way to go. It is interesting
>how open they all were in describing what they were controlling for.

It isn't really necessary to educate everyone about control theory. If HPCT is right, everyone has the necessary levels of control and perception. What's needed is a set of experiences in which they can see the basis for justice and constitutional rights. I can imagine role-playing in school situations where students are shown how ANYONE can be falsely accused and suffer prejudice. You don't have to be black or poor to be treated according to prejudices about a group to which you belong. This can happen to salespeople, policemen, rich people, Polish people, actors, garbagemen, psychologists, computer programmers, teachers, and students. Statistical treatment of test results formalizes this kind of prejudice, but it happens to everyone, or could.

What's important is to recognize that any person can suffer unjust accusation and bad treatment when they're identified as members of a population instead of being treated as individuals. The shared system concepts under which we agree to live have to be designed to work no matter which side of an accusation you find yourself on. If teaching this idea were made the core of courses in civics, maybe people wouldn't grow up to be jurors who don't get the point.

>Had these jurors been raised in an environment in which truth and
>justice were very important aspects of their System Concepts, their
>perceptions of the evidence would have been different.

AMEN.

Technically Sweet (920628) -- (what were your parents thinking of?)

Welcome to CSGnet, TS. As Rick Marken said, we're aware of Brooks' subsumption architecture, but he doesn't seem to think there's anything interesting for him in control theory.

An addendum to Rick's remarks. The main thing missing from Brooks' approach is an appreciation of what any control problem looks like from inside the controlling system. It's all very cute to build a robot that will go around collecting pop cans (although I understand that descriptions of this device's behavior are extremely overblown). But to do this, there has to be a designer who knows a hell of a lot about the properties of pop cans, environments, physics, and so on. When you see behavior as the accomplishment of a bunch of objective "tasks," you can give the behaving system a lot of help -- laser range-finders, for example. You can place its sensors so they are just right for responding to things you know are critical to the task, and build into its computers all the knowledge it will need about properties of its environment. Most of what makes these little robots at all successful resides outside the robots, in the PhDs of the designers.

If you put yourself inside a robot, everything disappears except what your sensors tell you. That's the only world there is. Suddenly there aren't any objects, or people, or motions, or even spatial attributes of reality. All you have are sensor signals. Whatever you accomplish has to be defined in terms of those signals, not in terms of their implications in a world that a human observer can experience. The effects of your outputs are known only in terms of the way the set of sensor signals changes. There's nobody to tell you what's really happening outside you when you generate an output. All you know of the consequences of acting is that your sensor signals change. If you're going to make any sense of anything, it can be done only in terms of the effects of your output signals on your input signals. What goes on out there between output and input is hidden from you.

The problem with the task-oriented approach is that it's all about what happens OUTSIDE the robot, as seen by a sophisticated human observer and not as seen by the robot. The human designer can set up the robot so that when it does things inside its own perceptual world, something of interest to the designer happens outside the robot's ken, as a side-effect. I think this is a misconception of what robot design is really about -- it's just a way for a human being to accomplish something simple, of significance to the human being, the hard way. Nothing has really been learned about the design of autonomous systems.

Perhaps you will agree with this: what Brooks and others following similar paths don't realize is that they themselves ALREADY live in a Virtual Reality (VR). They are inside a brain that knows only what its sensors tell it -- the outputs of all those sensors ARE the world, as far as they or anyone else can tell. They only way they know anything about the physical effects of their own actions (or anyone else's) is to examine their own perceptions. Failure to realize this is the reason behind their rejection of (or failure to grasp) the PCT principle of behavioral modeling. And it's the reason behind their task-oriented approach to modeling behaving systems.

As a person interested in VR, you probably see very clearly that behavior has to be organized around making perceptions behave in particular ways, not around "objective" effects of actions. When you set up a virtual reality and display it on a person's retinas, you're simply creating a different link between the person's actions and the person's (visual) perceptions. You can put any properties you like into this link, and eventually a person will learn to control the resulting perceptions by generating the required actions. It's only a tiny step to realizing that this is how natural behavior works, too. People know no more about the real physical links between actions and natural perceptions than they know about your computer programs for creating apparent links in a Virtual Reality.

By the way, the use of a computer running Lisp for controlling simple behaviors is overkill of the most extreme degree. A lot of very interesting behavior could be accomplished with a control system consisting of two or three operational amplifiers costing about 50 cents apiece. The world has gone digital and has apparently forgotten that there's a much easier (and far faster) way to accomplish analog tasks.

Best to all, Bill P.

Date: Mon Jun 29, 1992 9:19 am PST
From: Dag Forssell / MCI ID: 474-2580

Subject: Taylors diagram

[From Dag Forssell (920629-1)

Martin Taylor 920626 17:45

I, too, find Martin's diagram very helpful. I think I have understood for some time the complexities of "Behavior of perception," with control circuits in a suggested 11 levels and rather massive parallel. The insight that is at the same time simple and new is the view of the world outside the organism.

As I interpret the diagram, Martin is simply placing a mirror at the interface. Instead of thinking of the physical world outside as one complex of parallel happenings, I now realize that since "It is ALL perception," we must by necessity think of it as perceived by us with precisely the same levels of perception as we think of and discuss as inside the control system.

I think it is difficult to convey the significance of the rubber band experiment. Now it is easier.

With all 11 levels laid out both above (as always) and below (which is new) a mirroring line of symmetry (the border between the control system and the outside world (as perceived), we can portray the perception of knot and the dot as simple configurations in the outside world, controlled from the simple configuration want in the inside world.

This spans level +4 (above) past 0(interface) to -4 (below), or 9 levels in a 23 level diagram. (Perhaps better relationships: +6 0 -6)

Even the simple rubber band involves millions of sensors (on the retina, joint angles) and hundreds of effectors (muscle control) to develop and control the perception of a relationship between two single points in space.

Now you can visually (in the diagram) extrapolate and discuss how you control more complex matters like principles and even systems concepts involving even more sensors and effectors.

It becomes a straightforward extension involving, ultimately, all 23 levels, with a massive number of sensors and effectors at level 0.

Here we have separated perception and control and lost the sense of interconnection of the control processes up the levels, which is the essence of "behavior of perception," but this may be a better place to start. The simultaneous behavior of perception is a concept and process that is hard to grasp. Perhaps it may just as well wait a little.

It just occurs to me that we take Martins 23 level chart and fold it on the mirror line, then interconnect the control systems across so we control all the perceptions up and down. We are back to the diagram as we know it, but with an expanded understanding of it.

Best Dag

Date: Mon Jun 29, 1992 12:36 pm PST
Subject: Bicker about Bickhard

[from Joel Judd]

STOPPED SHAKING YET, RICK?

Bill, Rick, Martin, Bruce, Tom:

Gary might not have provided enough of his Bickhard conversation to provoke some discussion. Of course if you're controlling for not seeing more of the conversation then this post won't help. But I think he IS on the epistemological track, and that is something sorely needed. I am also personally interested in getting feedback on a question regarding the mechanics of an HPCT, so humor me for a bit.

On the epistemology, he is adamant that anyone dealing with knowledge systems recognize that when one is at the point of modelling LEARNING, then the system has to be able to develop interactions with its environment without knowing what is to be represented in those interactions. As soon as one assumes representation of knowledge in a system then one begins down the road of infinite regress. In education, this fallacy leads to the "transmission" notion of teaching: the teacher has some knowledge the student doesn't, and must somehow transmit that knowledge to the student. This is a paradox of learning that is yet to be widely admitted and studied in educational theory. Currently, the field of study that is distracting many from addressing this fallacy head-on is neural nets, because they APPEAR to be learning, but in fact the nets are not developing their own knowledge from scratch--they are deriving new combinations of prior knowledge (provided by a programmer).

So while we may be a ways from modelling knowing systems from scratch, I think that this basic epistemelogical fact needs always to be kept in mind--especially in an educational setting. In this part of the conversation Bickhard says:

"...in classic wax slate models errors are inevitable but the more you can avoid them the better. Well try applying that to any skill learning. You don't learn to drive, you don't learn to walk, you don't learn to play golf, you don't learn any skill by practicing all of the subelements to perfection and then combining them and pracitcing that to perfection. That is absolutely the worst conceivable way to learn anything. And yet at the cognitive level we assume that that's the optimal way to learn anything...people learn things by progressive approximations just like they learn any other skill and when they learn what this approximation needs to do in order to become a better approximation, in effect what they have done is to learn a new critical principle. They have learned a way that this approximation still makes some sort of an error or another and by learning in that manner they end up understanding what they've learned because they have learned all of the error criteria that it is satisfying that make it an acceptable or good or whatever proposal.."

At this point Gary mentions that just error doesn't give you the solution, which Bickhard acknowledges by pointing out the role of environmental influences and, in education, the classroom, and how good teachers are ones that appreciate the process of "scaffolding" or working one's way through problems. On a personal note, I don't feel I am as good a teacher as others because I tend to FORGET how I learn something, whereas good teachers (I think) tend to REMEMBER what it took to learn something and are adept at identifying and appreciating other who are going through similar processes.

The point here is that Bickhard's epistemological arguments are basically correct and there educational implications (as well as present or future implications for AI, robotics, etc.) should be recognized and remembered, even if the mechanics of his model don't interest PCTers.

And to the mechanics...

I think, I repeat I THINK, part of Bickhard's problem with reference levels has to do with their origins. So I'm going to go ahead and ask this question without consulting PCT scripture. If one understands a reference level to be a "goal" in the sense that the organism is acting to perceive perceptual inputs similar to the reference signal, then the reference signal can only come from two sources: memory of past experience (i.e., the recording of a past reference signal), or my imagination of a reference signal. In either case, in order to EXPLAIN the system, one would like to be able to explain where the reference signals come from. In the case of LEARNING, memory signals would seem in a sense to be trivial, since they result from prior learning (though they are anythign but trivial for, say, "practice" or "fluency" or whatever). Imagination signals would seem to be key in learning something new. But we're faced with the dilemma: How do I know what I am trying to learn? In other words, How do I have a reference signal for experience I have not had? How do I know when my reference signal for 'ride a bike' or 'graduate from the university' have been satisfied (I purposely chose these obvious examples)? More to my and some other's interests: How do I know when my reference signal for 'learn another language' is satisfied? I think in a considered answer to those questions lies much of the variability one observes in others' attempts to achieve their goals.

FInally (really) [is Lubin still out there?]: is there any neurophysiological support for the following supposition (and here in fairness Bickhard is just responding to Gary's query):

"...I argue that the nervous system (NS) should be looked at as--well, even this is only a first approximation--as a complexly organized system of oscillators that modulate each others' activity. Now that's only a first approximation and I think a closer one is the NS should be looked at as a complex topology of media for oscillations to move in. And those oscillations move in the medium in accordance with whatever the local properties of the medium are and in accordance with this topology and thereby modulate each other. That's the general architecture I would argue for. And one demonstration that something like that ought to be the case is the fact that a very large proportion of our CNS neurons are silent, they never fire...but in terms of modulatory effects they can be doing a lot." (Such as affecting the ionic concentrations in the area which affect dendritic computations, etc.)

Are topologies in use in neurophysiology? In engineering? Do they make sense in CT terms? Can someone provide a laymen's explanation of why they are or aren't to help my conceptualization of the brain?

Awaiting replies....

Date: Mon Jun 29, 1992 2:05 pm PST
Subject: Re: Bicker about Bickhard

[Martin Taylor 920629 17:00]
(Joel Judd, undated (920629?))

>I am also
>personally interested in getting feedback on a question regarding the
>mechanics of an HPCT, so humor me for a bit.

The questions that follow don't seem to need humoring.

> [Bickhard says that] anyone dealing with knowledge
>systems recognize that when one is at the point of modelling LEARNING, then
>the system has to be able to develop interactions with its environment
>without knowing what is to be represented in those interactions. As soon as
>one assumes representation of knowledge in a system then one begins down
>the road of infinite regress.

No, there's no infinite regress here, unless one's arguing style is singularly obtuse. Knowledge can certainly be represented, but interactions do not require the kind of representation that implies regress.

I think the fallacy here is the usual one perpetrated by philosophers, that of using a word in two different senses, leading to an apparent contradiction. Only from a cognitive-symbolic model does one come to the notion that there is a problem in "interacti[ng] with its environment without knowing what is to be represented in those interactions." "Knowing what is to be represented" implies that the actor (interactor?) is viewing its own performance from outside and conducting it by means of a set of rules based on some knowledge representation. Interactions performed by a control system don't need to work that way, and indeed will not work that way except at levels above programme. But there is a representation, based in HPCT on the reorganization that alters the sign of the gain function that turns an error signal into the reference for a lower-level control system.

There is another representation, in the perceptual combining function for each ECS. One may ask how that representation is learned, but it is not the same question as "how does the system develop interactions with the environment
without knowing what is to be represented in those interactions." The system as a whole does not "know" what is in these perceptual functions, but it can learn them. Simple Hebbian learning may develop them, or genetic algorithms
(as we intend to attempt), or something akin to reorganization. The point here is that if there is a complex environmental variable susceptible to control in the environment, and an ECS controls it, even poorly, then that ECS can adapt to control it better, by modification of its perceptual function.
No "knowledge" is required, unless one asserts that the perceptual function

that adapts is its own knowledge representation. There's certainly no infinite regress implied by that.

> In education, this fallacy leads to the
>"transmission" notion of teaching: the teacher has some knowledge the
>student doesn't, and must somehow transmit that knowledge to the student.
>This is a paradox of learning that is yet to be widely admitted and studied
>in educational theory.

There's no paradox that I can see. The teacher does indeed have knowledge that the student doesn't, and must somehow transmit it. The fallacy is in thinking that the teacher can do this if the student doesn't have the building blocks for the new knowledge. To build a house needs bricks, pipes, wire, paint, and myriads of other things, but just to pile them up doesn't make a house. You have to organize them in previously unknown ways (new knowledge). Neither can you make a house unless you have the bricks, etc.

In our BLC theory of reading (Taylor and Taylor, The Psychology of Reading, Acvademic Press 1983) we discussed what we called a "three-phase" pattern of learning: (1) acquisition of the perception of wholistic structures which are subdividable in some way, (2) the perception of the subdivisions and the relations that can and cannot occur among them, and (3) a new, more precise perception of the whole structures, which now includes the perception of the relationships among the units. Of these phases, only (2) can be "taught." (1) and (3) are "learned." In reading, (2) corresponds to the teaching of phonics, and it is the teacher's responsibility to provide the circumstances in which the regularities that permit subdivision can be perceived. A teacher can then inform the student about the ways the units can and cannot fit together. Of course, if you already have the units, new ways of organizing them can be described and "taught." So you can build both up and down from any suitable percept or set of percepts.

There was an experiment done in the late 50s or early 60s by Wilson P (Spike) Tanner at Michigan that illustrates the ability of people to develop brand new perceptual functions. His claim was that through feedback, people could learn to discriminate any pair of auditory signals that were physically distinct. He designed a pair of signals that nobody could tell apart, but that were physically quite different. Then he put people into a psychophysical discrimination study using those signals, for a long time each day for many days. Typically, the subject would get 50% correct (no discrimination) for many days, and would insist that the signals were identical. Then, one day, some distinction might be observed, and very quickly, perhaps within the same day, the score would rise to near 100% (perfect discrimination). One subject, as I recall, took 43 days before catching on.

>Currently, the field of study that is distracting
>many from addressing this fallacy head-on is neural nets, because they
>APPEAR to be learning, but in fact the nets are not developing their own
>knowledge from scratch--they are deriving new combinations of prior
>knowledge (provided by a programmer).

And these combinations ARE new knowledge. No-one develops their knowledge (if by that you mean performance ability) from scratch. We come from a few billion years of evolution, and have many building blocks available at (and before) birth. We develop new ones from the originals, and use them as

building blocks. Programmers seldom know what their neural networks are doing,
so if all the net's behaviour is knowledge provided by the programmer, it
is knowledge the programmer did not know she had. A strange kind of knowledge.

Does Bickhard claim that one knows only what one can teach?

Martin

Date: Mon Jun 29, 1992 2:14 pm PST
Subject: Re: Taylors diagram

[Martin Taylor 920629 1515]
(Dag Forssell 920629-1)

>It just occurs to me that we take Martin's 23 level chart and fold it on
>the mirror line, then interconnect the control systems across so we
>control all the perceptions up and down. We are back to the diagram as
>we know it, but with an expanded understanding of it.

This is true, but I'd rather not do that. One of the points of making the
diagram is to show a relation between a "Boss Reality" and a controlled
percept. Now if that Boss Reality exists, it is accessible to another
control system (e.g. an experimenter). The other control system can
focus on (perceive) the same complex environmental variable (CEV) and perhaps
attempt to control its perception of the CEV, disturbing the first control
system's perception of it. We can diagram the interaction his way, taking
the top half of this diagram as representing the whole of my earlier one, and
the bottom half its mirror image:

```
                ECS
  Person 1     /   \
              /|   |\
         ====^=^===V=V====
              \|   |/
               \   /
   World        CEV
               /   \
              /|   |\
        =====V=V===^=^======
              \|   |/
  Person 2     \   /
                ECS
```

If you fold the original diagram about the mirror line where "the rubber meets
to road", you lose this view. Much more crucial, you lose the view of somewhat
off-focus disturbance, which was the core of the discussion about VOT.

In the second diagram (below), Person 2 is attempting to disturb what he/she
thinks Person 1 is controlling for, but actually disturbs only one of Person 1's
intermediate percepts. Person 1 might not be controlling that percept with

high gain, but might control the focus CEV by compensating for error in the disturbed intermediate variable through a shift in the reference for another intermediate variable. This would be detected by Person 2 as variability or lack of control in the disturbed intermediate variable, and Person 2 would miss the fact that the focus CEV was under tight high-gain control.

```
                ECS
  Person 1     /   \
               |   ECS
              /|   |\
         ====^=^===V=V====
              \|   \ |
               |    CEV\
               \   /  | \
   World        CEV  /| |\
                    / | | \
       ============^==^=V==V======
                    \ | | /
    Person 2         \| |/
                      \ /
                      ECS
```

I would have preferred to draw this diagram with the Person 2 picture aligned at right-angles to the Person 1 picture, and if I do it with a graphics system,
that is what I will do. The right-angles presentation shows more clearly the missed aim of Person 2. But it is hard to do in ASCII.

Martin

Date: Mon Jun 29, 1992 3:17 pm PST
Subject: Mirror symmetry; Meno paradox

[From Bill Powers (920629.1600)]

Dag Forssell (920629) --

Thanks for the 23-level system with mirror symmetry about level 0. The 11 levels below the line through level 0 are, of course, in the observer.

We communicate with and study other control systems through level 0. Physics, supposedly, studies level 0 in itself -- but it inevitably contains elements of levels 1 to 11 in the observer. Hence "particles," "forces," "space," and mathematical physics. The big unsolved question is the extent to which the 11 levels in the observer are shaped by the existence of organization in level 0.

Joel Judd (920629) --

RE: Bickard

>On the epistemology, he is adamant that anyone dealing with knowledge
>systems recognize that [w]hen one is at the point of modelling
>LEARNING, then the system has to be able to develop interactions with

>its environment without knowing what is to be represented in those
>interactions.

This is the Meno paradox that Hugh Petrie wrote about:

"You argue," say Socrates, "that a man cannot enquire either about that which he knows or about that which he does not know; for if he knows, he has no need to enquire, and if not, he cannot; for he does not know the very subject about which he is to enquire."

This "paradox" probably arises from stating the problem badly. I remember having a similar problem when I was learning about computers. If you know what the data in memory is, you don't need addresses to find it. But if you don't know the data, you don't know how to find its address in memory. I went around and around several times before I figured out what the problem was.

Another: I know that there are light switches and I know that there are lights in the room. What I don't know is which switch is connected to which light. Clearly, I can't know which switch to use to turn on a given light until I know what the connections are, but I can't reason out the connections until I know which switch turns on which light. Like Zeno, I conclude that the simple is impossible. And like the answer to Zeno's Paradox, the solution lies in stating the problem differently.

What actually happens is that I flip a switch WITHOUT knowing what the connections are, and thereby discover them. We find new knowledge by making mistakes -- guessing wrong. Socrates and Meno were trying to solve the problem without trying any actions. They thought that the solution could be arrived at by pure reason, by passive consideration of perceptions, by logic. But the actual solution requires putting logic aside and trying something at random.

In the case of computer memory, what was missing was the fact that some programmer ASSIGNS meaning to memory addresses arbitrarily, without knowing what values will be stored in those addressed locations. Sorting this out requires learning the difference between the NAME of a piece of data and the VALUE of that piece of data -- it's exactly the same problem that confronts students when they start to learn algebra. "If x doesn't have any value, how can you say it's equal to y?" Wrong question.

So what about new knowledge in general? If you don't already perceive a collection of sensations as an object, how can you ever learn to perceive that it's a square? I think the answer is that you try many different ways of perceiving, and by experimenting with actions, learn which perceptions are controllable and which aren't. Then you name them.

In my model there are predispositions to perceive in 11 different styles, hierarchically related. In principle, these predospositions could have arisen through evolution because these are the most useful categories in which to perceive and control -- useful in terms of staying alive long enough to reproduce and enhance the reproduction of your conspecifics (the latter being the Bawd Master theory of natural selection). This would explain why one eventually perceives order in a collection of sensations by creating configuration perceptions instead of some other kind. But the question of why we learn to perceive only certain specific configurations pertinent in the current environment can be explained only in terms of a-logical reorganization. Reorganization is the basic solution of the Meno paradox.

Part of Bickhard's problem is using the word "learning" without analyzing the phenomenon. No reorganization at all is needed to learn your telephone number: all I have to do is ask and memorize your answer. No reorganization is needed to learn the values of x and y in the pair of equations 2y = 3x + 4 and 5y = 6x + 1. All I have to do is apply a known algorithm, and when it is done I will have learned the answer. The only kind of learning that's at all a mystery is the kind that can't be done by relying on either memory or a fixed procedure. That's the kind of learning that requires reorganization.

With regard to new reference levels, I think it's clearly part of the HPCT model to say that reference levels can't exist before perceptions exist: all a reference level is is a specified value of a particular existing perception, specified by a reference signal having that value. So once you've settled on an answer to the question of where new perceptions come from, you've automatically answered the question of where new reference signals come from. Imagining a new reference signal doesn't mean imagining a new KIND of reference signal. It just means imagining a value for an existing kind. The answer to how you have a new reference signal for something you've never experienced is (a) you can't, if you mean a KIND of experience, or (b) easily, if all you mean is a new magnitude of a familiar experience.

You can, of course, imagine new combinations of values of familiar experiences. But these new combinations will have no significance to you until you've developed a higher-level perceptual function that recognizes the COMBINATION as something different from the INDIVIDUAL EXPERIENCES. After that happens, of course, and after you've acquired control of the new perception, you can choose new values for that perception to serve as reference levels.

You quote Bickhard as saying "...I argue that the nervous system (NS) should be looked at as--well, even this is only a first approximation--as a complexly organized system of oscillators that modulate each others' activity."

Why oscillators? The only reason I can think of is that chaotic systems are commonly characterized as oscillators that can go into different nonlinear modes. There's nothing about any behavioral systems I know about that suggests oscillatory phenomena except circadian rhythms and things like walking and keeping time -- which are striking because of their regularity, not their chaotic-ness. Birkhard is just looking for a place to stick chaos into his model along with all the other trendy arm-waving stuff like "modulation" and "topology." Again, from your quote of Birkhard:

>And those oscillations move in the medium in accordance with whatever >the local properties of the medium are and in accordance with this >topology and thereby modulate each other. That's the general >architecture I would argue for.

Whatever mental picture there is in Birkhard's head that goes with these words, it isn't ready for publication yet. It isn't even an architecture. Rick Marken calls this "high-falutin' incoherence." I agree. Birkhard is just getting off on his own brilliance. That's how the Scholastics used to make a living.

Birkhard is playing a game that many recognize and find irresistible. I recognize it, but find it easy to resist and even to loathe.

Best to all, Bill P.

Date: Mon Jun 29, 1992 3:51 pm PST
Subject: causation and misc

In response to Bill (920626)
Joint determination, Hmmmm. Well, I certainly agree that the reference signal plus all the other influences on the controlled variable completely account for the action. I would say the same for the outcome, too. I have no problem saying that this means that all of these things determine the action (or outcome) either. But my purpose in using these terms is to note the relationship between disturbances and actions (an inverse relationship in your tracking tasks) and referenced perceptions and outcomes (a strong positive relationship in your tracking tasks). Because if we define determine as "A determines B, if, given A, B is completely predictible..." then it seems my use of the terms is true to the phenomenon. I easily embrace what you said about joint determination, although it seems to imply that prediction is impossible since one cannot know all the influences. I am saying that prediction is possible without knowing all the variables and I would think you would say the same.

Thanks for clearing up the issue on temporality and causality. The PV=nrt example was used last semester in a grad philosophy class as a nonexample of causality. I suppose it there had been a physics student in the class we would not have come up with this conclusion. Your explanation makes alot more sense to me--the balloon example was always disatisfying even after we "understood" it.

Martin Taylor (920626)
Perhaps I missed a post, but could you help me understand the differences between physical observables, low-level interpretation, more-abstract interpretation, and a "solid" testable interpretation, and the remaining terms you use on the "World" side of your diagram? The idea of levels in the world is odd to me. Is the LINE a mirror, for I do not understand why the World items become more abstact as they move AWAY from the organism. I think I am missing something simple, but I can't get past my way of perceiving the diagram.

Carpe' Diem Mark

Date: Tue Jun 30, 1992 1:48 am PST
Subject: Re: Teaching system concepts; Brooks & PCT

[From Oded Maler 920630]

[From Bill Powers (920629.0930)]
..
>Perhaps you will agree with this: what Brooks and others following similar,
>paths don't realize is that they themselves ALREADY live in a Virtual
>Reality (VR). They are inside a brain that knows only what its sensors tell
>it -- the outputs of all those sensors ARE the world, as far as they or
>anyone else can tell. They only way they know anything about the physical
>effects of their own actions (or anyone else's) is to examine their own
>perceptions. Failure to realize this is the reason behind their rejection
>of (or failure to grasp) the PCT principle of behavioral modeling. And it's
>the reason behind their task-oriented approach to modeling behaving
>systems.

I think you misinterpret Brooks a little bit. As a roboticist he is surely aware that the robot's world is just the perceptual space of its sensors. I think all his work in the last years drove AI-based robotics much closer to the lines of HPCT. I *really* recommend reading his "Intelligence without Reason" [IJCAI91, available thru ftp] in order to see the context. (Of course, in the modern science game, hype is inevitable, but such things exist in other forms in other small communities not belonging to the scientific establishment..).

In particular, in section 6.2 when he describes the robot that picks
the soda can he says:

".. The hand has a grasp reflex that operated whenever something broke an infra-red beam between the fingers. When the arm located a soda can with its local sensors, it simply drove the hand so that the two fingers lined up on on either side of the can. The hand then independently grasped the can. ... As soon as it was grasped, the arm retracted - it didn't matter whether it was a soda can that was intensionally grasped or one that magically appeared."

[Please don't pick on the word "intensionally" in the last sentence, it just means that the soda can which initially "caused" the hand to approach the table, was not necessarily the one whose perceived grasping "caused" the hand to retract.]

[[And don't pick on the word "cause", etc.]]

I think an experienced reader can rephrase all this in terms of reference signals and controlled perceptions if he/she/it finds it useful.

Best regards

--Oded

Date: Tue Jun 30, 1992 6:15 am PST
Subject: learning

Could someone offer a HPCT definition of learning. I know that learning occurs via reorganization, but that doesn't tell me what it is. Is it a permanent change in reference values? Is it the creation of a new reference signal? A new comparator? Are there different forms of learning? It seems that I learn HOW to do something (ride a bike) and it seems that I also learn WHEN to do something (when I feel like x, I should stay home (or go out) OR when I can't resolve something, try not thinking about it for a while). I wouldn't be suprised if HOW and WHEN are different ways of experiencing the same thing, but they may not be. Either way, my questions remain.

Has anyone ever suggested a hierarchical reorganization system? If I was God, I'd prefer that there wasn't such a thing, but what about "learning to learn"--what's the HPCT explanation for that? And for clarification, in my mind I visualize the reorganization system "perpendicular" to the hierarchy, wherein the "input" is the error signal, and the "output" feeds into the reference signal (or perhaps the "output" of the system above), closing the loop. The "disturbance" is the perceptual signal. Am I close?

Mark

Date: Tue Jun 30, 1992 6:51 am PST
Subject: Re: causation and misc

[Martin Taylor 920630 10:10]
(Mark Olson undated ?920629)

>Martin Taylor (920626)
>Perhaps I missed a post, but could you help me understand the differences
>between physical observables, low-level interpretation, more-abstract
>interpretation, and a "solid" testable interpretation, and the remaining
>terms you use on the "World" side of your diagram? The idea of levels in
>the world is odd to me. Is the LINE a mirror, for I do not understand why
>the World items become more abstact as they move AWAY from the organism. I
>think I am missing something simple, but I can't get past my way of
>perceiving the diagram.

What the World IS depends on how we perceive it. At present, physics tells us it is a soup of interacting quarks, gluons, and so forth. Ordinary people don't perceive it that way. We see chairs, friends, books; we perceive happiness, group cohesions, or whatever. The interface between the control sytem that is a person and the world outside is very public and generally observable and measurable. It consists of photons, pressure waves, and the like (if we accept the entities that are currently the currency of physics). To perceive something useful, like a chair, the incoming data (fluctuations in photon flus, etc.) must be transformed through many layers of abstraction. The chair itself is an abstraction, an arbitrary separation of the quarks of the world into those that form part of the chair and those that don't. It exists only in the mind of the perceiver (and if we believe in evolutionary efficiency it exists only if it can be a controlled perception).

In HPCT, control of high-level perceptions is exercised through the mediation of control of lower level perceptions. A high-level perception is very abstract. So, therefore, is the complex environmental variable (CEV) it controls. The intermediary controlled percepts, and hence the corresponding CEVs, are less abstract. Democracy is more abstract than voting, which is more abstract than putting a piece of paper in a box, which is more abstract than pressing a pencil onto the paper...

But from the point of view of the ECS that controls a percept, the CEV that it is perceiving is the one solid and testable aspect of the world. It is also the most abstract entity involved in the whole hierarchy that participates in its control.

As I see it, a CEV and the corresponding percept in an ECS are not so much mirror images as corresponding foci. What lies in between is distributed and hard to identify uniquely with any aspect of the controlled percept or the CEV that leads to the percept. If you want another (imperfect) metaphor, a hologram of an object derives from a well-defined scene, and can recreate a convincing image of that scene, but no part of the hologram can be credited with responsibility for recreating any specific element of the scene. The hologram is concrete and physical, the recreated scene is abstract, though it may look solid.

If you missed a posting, it was one of Bill's, which I thought could be illustrated by my diagram.

Martin

Date: Tue Jun 30, 1992 7:55 am PST
Subject: Subsumption vs. HPCT

[From Bill Powers (920630.0800)]

Oded Maler (920630) --

>I think you misinterpret Brooks a little bit. As a roboticist he is
>surely aware that the robot's world is just the perceptual space of
>its sensors. I think all his work in the last years drove AI-based
>robotics much closer to the lines of HPCT.

I have a copy of Rodney Brooks' "Intelligence Without Reason" -- in fact several, thanks to my friends. The one I'm looking at is marked "To appear: IJCAI-91," so I've been aware of Brooks's work for some time.

I agree with many things Brooks says about AI, and with parts of his alternative to it. I have agreed with some aspects of his approach since I began this work nearly 40 years ago -- for example, behavior based on real-time interaction with the world, the emergence of complex behavior from simple interactions, the concept of parallelism, the idea of layering (although different from his idea), the building up of complex behaviors out of simple behaving units. I even agree with his assessment of cybernetics, especially its "lack of mechanism for abstractly describing behavior at a level below the complete behavior, so that an implementation could reflect those simpler components."

But Brook knows nothing about my work, or at least dismisses it (he never cites it). My model has employed situatedness, embodiment, the dynamics of interaction with the world, and emergence of intelligent-looking behavior from interaction of components of the system, right from the beginning. Brooks called these "a set of key ideas that have led to a new style of Artificial Intelligence research which we are calling behavior-based robots." If I were the one with a large budget and an academic position, I would be tempted to dismiss Brooks as having reinvented a (somewhat square) wheel. The main reason I can't do this (other than my natural sunny temperament) is that I have never had any support for building robots for fun; while I've long said that the environment is its own best simulation, actually using it that way, by building behaving robots, has been beyond my means. Of course I HAVE built robots, but nobody would recognize them as such.

You have to realize how work like Brooks' looks to me. I understand that from other points of view, it might seem that I would be encouraged to see that my ideas are similar to those of a leader in today's explorations of simulated behavior. From my point of view, however, Brooks is just starting to catch on to an approach with which I have had decades of experience. He still carries a lot of the baggage of old-style AI with him, despite his revolt against it. He pays little attention to real human and animal behavior, and so has missed most of the hints which it contains. He has not even discovered the simplest principle of control behavior, which is that the perceptual signal, not the act, is under control.

A year or two ago, after seeing some of the PR material on subsumption architecture in Science News and Science, I sent Brooks a long letter inviting him to cooperate with the CSG and challenging him to a discussion of issues. I sent him my Demo 1 and Demo2 programs, and version 1 of the Little Man. He did not object to anything I said or to anything in these programs -- at least not to my knowledge, as he has never replied. Unless he simply dropped everything in the wastebasket without looking at it, he evidently found nothing of any interest in the letter or the programs. I have given up pounding on closed doors.

Brooks lists under "Vistas" some "key research areas that need to be addressed ..." The first is "Understanding the dynamics of how an individual behavior couples with the environment via the robot's sensors and actuators." In other words he's about to start looking into the phenomena that will lead to re-inventing control theory, real soon now. He could save himself a lot of time and trouble if he were interested in seeing how others have solved his problem (starting some 50 years ago). But I think he's not much interested in going about it that way. He's convinced that he's doing something new and is way out ahead of the world. Well, that's far from true.

>In particular, in section 6.2 when he describes the robot that picks
>the soda can he says:
>
>".. The hand has a grasp reflex that operated whenever something broke
>an infra-red beam between the fingers. When the arm located a soda can
>with its local sensors, it simply drove the hand so that the two >fingers
lined up on on either side of the can. The hand then >independently grasped
the can. ... As soon as it was grasped, the arm >retracted - it didn't
matter whether it was a soda can that was >intensionally grasped or one
that magically appeared."

In fact, it probably didn't matter whether it was a soda can or someone's leg or a bug alighting on the infrared sensor. This is what I mean about the designer putting too much of his intelligence into the robot -- or at least into interpretations of what the robot is doing. This robot is billed as "reliably picking up soda cans." The casual reader gets the idea that this robot knows what a soda can is. It doesn't. Brooks implies that the array of scanners was sufficient to identify a "soda-can-like object." If that's so, he has achieved an enormous breakthrough in object recognition. I rather suspect that all these soda cans were upright, located so they couldn't be confused with a small book standing on a table, placed within easy reach, contrasting strongly with the background, and set up so they couldn't be knocked over too easily. I also doubt that when the hand grasped the can, it could tell what object it was holding, if any. I'm just imagining all that, of course.

"By carefully designing the rules," says Brooks, "Herbert was guaranteed, with reasonably reliability, to retrace its path." But Herbert did not know how to get back where it started; the rules said nothing about that. They said things like "When passing through a door southbound, turn left." There's the modeler's intelligence USING the robot to achieve an end that the robot itself is helpless to achieve. This is what I mean by hanging onto old-style AI concepts.

My basic objection to the subsumption architecture is that it is enormously wasteful of resources. Every new module starts from scratch, with raw sensory data. When one module finishes its task and another sees an opportunity to act, the new module actually has to turn off the old one and take over completely. An "approach" module is turned off, and a "retract" module takes over, because if they both worked at the same time they would come into conflict. Why not a single "position" module, given a reference signal adjustable over the whole range from approach through stasis to withdrawal, using the same sensor and effector connections? Brooks has broken behavior down into TASKS, with modules designed to perform each task as an achievement in the objective world. Once any task has been accomplished, the module assigned to it becomes useless -- to do a different task, one has to design a new module from the ground up.

In the HPCT model, the levels of control are used by higher levels; the higher levels use not only the control capacities of the lower systems, but in many cases the perceptual interpretations developed up to that level. While each control system is specialized to control just one specific variable, it is a general-purpose control system: it can be used in any situation where that variable needs to be controlled, and control can be set to occur around any reference level. Control is not organized around objective "tasks," but around control of specific sensory variables, independently of what external task is being performed. While the present "pandemonium" concept does waste resources in a different way, this waste may be unavoidable (in any model). But there is never any need, in the HPCT model, to create a new perceptual function or a new output function that is an exact duplicate of an existing one.

Once the first level of control in the HPCT model is reasonably complete, it can be used by all subsequent levels in the performance of any behavior of any kind for any purpose. In fact, it constitutes the ONLY MEANS by which higher-level systems can act. Higher systems have no direct access to sensors or actuators. This is the general case; once any level of control-system exists, it is the ONLY MEANS by which higher levels can control their own perceptions.

The subsumption architecture, therefore, contrary to Technically Sweet's comment, is fundamentally different from the HPCT architecture -- far more complex and entailing far more duplication of function. What hierarchical relationships do exist consist mainly of turning one module off and another one on. The system doing this switching would use, in the HPCT model, very high-level functions in which it is all too easy to insert the experimenter's knowledge of the world without giving that knowledge to the model itself.

The HPCT model is oriented toward the control of variables in hierarchical relationships, with objective effects of doing so being side-effects of no concern to the model. The subsumption architecture is oriented toward producing certain objective effects and relationships as seen from the standpoint of the observer outside the robot; the robot is given whatever instructions and task-achieving modules are needed to make an effect occur in the perceptions of the observer. The HPCT model relies exclusively on feedback control of perception. The subsumption model uses feedback in only the crudest on-off way, and in many places employs the "SMPA" (sense-model-plan-act) principle that Brooks calls a "bias" that has impeded AI research. Brooks himself has characterized his robots as "stimulus-response" devices. And indeed, for the most part, that is what they are. In some places they accidentally incorporate control-system principles. But there is no principled application of control theory.

Brooks has a long way to go to catch up with HPCT, although others have a longer way to go.

Best, Bill P.

Date: Tue Jun 30, 1992 9:20 am PST
Subject: Re: Teaching system concepts; Brooks & PCT

[From Rick Marken (920630.0900)]

Oded Maler (920630) says:

>I think you misinterpret Brooks a little bit. As a roboticist he is
>surely aware that the robot's world is just the perceptual space of
>its sensors. I think all his work in the last years drove AI-based
>robotics much closer to the lines of HPCT.

Funny how psychologists, biologists, roboticists, etc have been "moving toward" HPCT for years. Randall Beer was supposedly close to PCT. Brooks is claimed to be close. All kinds of psychologists are pointed to as having PCT-like theories. Why don't they just adopt PCT, period. There is a reason -- it's because "close" is VERY FAR same "getting" PCT and getting PCT means kissing all the cause-effect crap goodby and recognizing that behavior is the control of perception. The difference between "moving toward" PCT and getting PCT is like the difference between "moving toward" being pregnant and being pregnant. Even the stupid, confused Supreme Court can apparently tell the difference between those two states. I just think Brooks, Beer, Carver and Scheier, Lord, and all the other people who are supposedly "moving toward" PCT just don't want to get knocked up (if they came to a CSG meeting we could arrange a nice marriage ceremony for them, however).

Best regards Rick

Date: Tue Jun 30, 1992 9:49 am PST
Subject: Learning; Mirror diagram

[From Bill Powers (920630.1100)]

Mark Olson (920630) --

>Could someone offer a HPCT definition of learning. I know that >learning occurs via reorganization, but that doesn't tell me what it >is. Is it a permanent change in reference values? Is it the creation >of a new reference signal? A new comparator? Are there different >forms of learning? It seems that I learn HOW to do something (ride a >bike) and it seems that I also learn WHEN to do something (when I feel >like x, I should stay home (or go out) OR when I can't resolve >something, try not thinking about it for a while). I wouldn't be >suprised if HOW and WHEN are different ways of experiencing the same >thing, but they may not be. Either way, my questions remain.

A couple of days ago I talked about "learning" as a fuzzy word with several unrelated meanings. You can learn an address by memorizing it. You can learn a method for solving equations by following a prodecure in a mathematics manual. Reorganization is concerned with learning for which there can be no rational basis and that doesn't depend just on experiencing and remembering something new.

Learning can't result in a "permanent reference level change" because reference levels are adjustable -- and must be -- in the HPCT model. Only the highest reference level, for system concepts, would tend to remain the same for long periods of time. All others change as higher-level systems encounter errors and try to correct them. Remember that adjusting a particular reference level has the effect of specifying the AMOUNT of a particular perception that is to be brought about and maintained.

It would help if you were to translate from informal language, like HOW and WHEN, into more specific controlled perceptions in terms of the model. Knowing HOW to ride a bicycle means learning what variables to control, and also acquiring the detailed input and output functions needed to monitor and affect the variables in a stable way. Reorganization is required if you've never done it before. You have to pick out many kinesthetic and visual variables, and keep changing the way you perceive until you're perceiving the right things in the right relationships. You have to alter the amount of output generated by a given amount of error, and also adjust the response to the first and second derivatives of the error, to keep from wobbling and losing control. When I say "you" have to, I mean that reorganization must have these effects -- you don't have much conscious control over that process, if any. All you can do consciously is keep trying and falling off.

Part of reorganization is an experimental shuffling of connections between levels of control, so that a higher-level error comes to be connected to the right lower-level reference signals. On the perceptual side, it's a shuffling of connections from lower-level perceptions to higher-level input functions, so you find ways of constructing higher-level perceptions out of particular lower-level ones. And there's also the matter of the FORM of the perceptual function -- how the higher-level perception will depend on the lower-level signals. That's produced by reorganization, too.

Learning, real learning that occurs through reorganization, is basically a random process. It isn't driven by what needs to be learned, but by the consequences of NOT learning SOMETHING that works. What "works" is defined not in terms of the perceptions and actions in the hierarchy of control, but by intrinsic error: deviations of critical variables from their inherited reference levels. As reorganization goes on, the form of behavior changes; perceptions change in relation to the environment, selection of lower-level goals changes, means of acting changes. If the changes result in bigger or the same intrinsic error, reorganization occurs again right away. If the change reduces intrinsic error, the rate of reorganization slows. When intrinsic error drops below some minimum amount, reorganization stops, and whatever organization of behavior exists at that moment has been "learned." It doesn't matter what that organization is, as long as it results in correction of intrinsic error.
There is no advance specification of what is to be learned; anything will do, as long as it has the side-effect of reducing intrinsic error.

So the reorganizing system basically doesn't care about behavior, perception, cognition, or any of those things we associate with conscious experience. With respect to the learned hierarchy, the reorganizing system acts like an unsympathetic boss or a cat: "I don't care how you keep me fed -- but whatever you're doing now isn't good enough, so change it!" But this boss, or cat, doesn't wait for you do change something. It reaches in and twiddles something whether you want it to or not, whether you like the result or not. And if it's still not satisfied, it does it again and again, relentlessly, until it IS satisfied.

I assume that the reorganizing system can make arbitrary changes in any part of the hierarchy from bottom to top. I assume that the changes are all quite small, and that they occur at a low enough rate that the consequences of each change have time to be reflected in the state of the intrinsic error signal. I assume that a law of small effects is at work: small changes have small effects. And I assume that there are some aspects of organization that are NOT changeable: the basic kit of neurons available at each level of control, for example, with properties that favor development of perceptual systems and output functions peculiar to that type of controlled variable.

>Has anyone ever suggested a hierarchical reorganization system?

Yes, the thought has crossed my mind. It's possible that one level of reorganization exists at the level of DNA; that another is involved in development from zygote to adult; that a third is involved with the biochemical systems (the immune systems, for example), and that a fourth is involved primarily with the learning of motor behavior and higher levels of neural organization.

>... in my mind I visualize the reorganization system "perpendicular" to >the hierarchy, wherein the "input" is the error signal, and the >"output" feeds into the reference signal (or perhaps the "output" of >the system above), closing the loop. The "disturbance" is the >perceptual signal. Am I close?

Well, not very. The reorganizing system, as I visualize it, is not concerned with the same variables that the neural hierarchy deals with. Think of the reorganizing system's effect not as depending on signals flowing in the neural hierarchy, but on physical states of the organism. Think of its actions as actually ALTERING THE WEIGHTS OF SYNAPSES AND EVEN THE CONNECTIONS FROM ONE NEURON TO ANOTHER. The reorganizing system is not concerned with the signals flowing in those pathways (or with their meanings), with the possible exception of error signals (one good reason for supposing that comparators are part of the built-in functions). It's only concerned with the consequences of particular ways of behaving on those variables it cares about: the ones on which life and continued existence depend, which are largely outside the purview of the learned systems.

If you look at the chapter on Learning in BCP you'll find some diagrams that may help, and further discussion.

Martin Taylor (20630.1010) --

Better and better. Something seems to be happening here after our long shakedown cruise. A very coherent picture is emerging. Maybe it's just that we have finally got past the barriers of language and have eliminated most of the irrelevant or illusory disagreements. At any rate, I'm enjoying it.

Best, Bill P.

Date: Tue Jun 30, 1992 11:28 am PST
Subject: On Vacation 920702-15

[from Gary A. Cziko 920630]

I will be in France and Switzerland from July 2 to July 15 and therefore unable to respond to e-mail during this time.

If anyone wants to stop receiving CSGnet traffic while he or she is away on vacation, send the following command to LISTSERV@VMD.CSO.UIUC.EDU as the first line in the text of an e-mail message

set csg-l nomail

And to resume receiving CSGnet traffic, send

set csg-l mail

To receive a log file of mail you have missed send the following command to LISTSERV

get csg-l log9207a

This will give you a file of all posts for the first week of July with log9207b the second week, etc.

I may have a few comments to make the net before taking off, but I wanted to get this out now before the last minute.--Gary

P.S. Why don't you all you CSGnetters out there spend the next two weeks thinking instead of bashing at the keys so that I don't have 400 message to wade through when I come back. Perhaps I should start the policy that the network is on vacation whenever the network manager is?

Gary A. Cziko

Date: Tue Jun 30, 1992 1:33 pm PST
Subject: Re: Learning; Mirror diagram

[Martin Taylor 920630 15:00]
(Bill Powers to Mark Olson 920630.1100)

>Part of reorganization is an experimental shuffling of connections between
>levels of control, so that a higher-level error comes to be connected to
>the right lower-level reference signals. On the perceptual side, it's a
>shuffling of connections from lower-level perceptions to higher-level input
>functions, so you find ways of constructing higher-level perceptions out of
>particular lower-level ones. And there's also the matter of the FORM of the
>perceptual function -- how the higher-level perception will depend on the
>lower-level signals. That's produced by reorganization, too.

I'm not sure I follow this. You seem to be talking about three different things under one label, "reorganization," and none of them are what I have hitherto thought of as reorganization.

So let's try to iron out another wrinkle. I'll try not to use the word at issue, atleast until after I list what I see as independent ways in which the hierarchic control system can improve its ability to control its percepts.

First, bear with me in what seems, but is not, a digression.

An issue that's been gnawing away for a while: I take the highest level references to be the ones you call "intrinsic." For the most part, they provide references for perceptions based on body chemistry, I assume. You keep talking of the System Level as the highest level that [controls?/provides references?]. If the System Level is the highest set of ECSs, where do the references for it come from? If the System Level simply sets reference levels for the highest level of ECSs, then where do the settings for these references come from? I have assumed that the intrinsic references are the references for percepts at the System Level. The alternative seems to be that the top level is not a control level at all, but simply a (possibly chaotic) dynamic system.

If the intrinsic references are indeed the top-level references, then the question of learning can be posed in two contexts: in evolution and in individual development. In each case, the organism starts with simple physical/chemical sensors and some way of controlling what they sense--the bacterium wiggles faster if things are wrong than if they are right, and thereby tends to get out of a harsh environment or stay in a benign one; perhaps the embryo also has some way of stabilizing its internal chemistry in a reasonable state--individual cells do. So, initially, we have a one-level control hierarchy, into which new ECSs must be inserted by evolution or by maturation to create new levels.

The "learning" question is how new ECSs get inserted at levels between the top and bottom, between the intrinsic references and the sensors/effectors. If the cleanly layered HPCT model is correct, humans have evolved so that they can "learn" ECSs at eleven different levels of a hierarchy. Other species may have fewer levels of ECSs, or may have ECSs with perceptual functions of quite different kinds (though if we look at the retention of other characteristics through evolution, it would be unlikely for us to have ECSs of kinds very different from those of other mammals.) But a 3-level control system would be expected to maintain control of its intrinsic percepts (i.e. percepts corresponding to its intrinsic references) much more crudely and over a narrower range of disturbances than would a 6-level control system.

All of the above is intended to argue that a major cause of change in a control hierarchy (learning) is the insertion of ECSs between existing ECSs and their input/output connections. Control becomes more subtle.

(Aside--this concept mirrors, perhaps not by accident, the notion of three-phased learning that we proposed in 1983 for reading skills.)

(Second Aside--there is a problem here if we maintain that there are no intra-level linkages such as "configurations of configurations." In the three-phase learning, what is centrally learned is that large configurations can usefully be divided into sets of related smaller ones, many of which occur in all sorts of larger ones; the letter "c" occurs in "cat" "occur" "sick". I take it that letters and words are both configurations. Experimentally, there is evidence that words can be perceived both directly and through their letters. If this is true, there is either level-jumping or intra-level connection, neither of which is consistent with the clean layered hierarchy. I don't want to pursue this argument further either here or in the immediate future, but I raise a flag to mark a potential topic for later consideration.)

Back to learning. Here are some mechanisms: (1) insert an ECS between two (or more) existing ECSs, with or without disruption of the existing links. (Third Aside: I don't see how one can get an ECS into the hierarchy in any other way. The new one must have a source for its reference, must get its perceptual input from somewhere, and must send its output somewhere. All those somewheres must be ECSs that already exist); (2) break links or make new links between existing ECSs; (3) change the sign on an output link of an ECS; (4) change some parameter of the perceptual input function of an ECS. There may be other learning mechanisms, but these four seem distinctly different, and within each there are several possibilities for further subdivision according to the conditions that allow them to happen.

Until I read Bill's posting, I took "reorganization" to be (3), applied in earlier discussions randomly throughout the hierarchy when the intrinsic error is high, but as Bill agreed some time ago, reorganization must happen in a much more modular fashion. There are too many degrees of freedom in the system to allow global reorganization to have much probability of improving control. But Bill seems to incorporate all of (2), (3) and (4) under the same name. I think they should have different names, and I propose "restructuring" for (2), "reorganization" for (3), and "adaptation" for (4).

What do these types of change in the network correspond to, in everyday language? (1) seems to come closest to maturation--"reading readiness" and the like, but it could include other possibilities that occur whenever we learn to perceive a new kind of complex environmental variable (CEV). I don't know where the new ECS would come from. There is no problem about its being a new kind of ECS, because "kind" is in the mind of the analyst. Except for the exact nature of the perceptual input function, all ECSs are more or less the same, regardless of what level of abstraction they actually control.

Restructuring--type (2) learning--can happen randomly, as can reorganization (type 3). But I would expect the occurrence to be strictly local. An ECS that was not maintaining control despite a high output gain would be likely to reorganize (reverse some random component(s) of its output without changing the nature of the actions it affects) or to restructure (randomly try to do something new). I see no rationale for either reorganization or restructuring to affect ECSs that are happily maintaining control, or that are operating with
a low gain (insistence). There should be no need for a "reorganizing system," a concept Bill has mentioned from time to time, including this posting. If a high-level ECS reorganizes or restructures, it changes the references it supplies to lower level ECSs (though it has no information that it does so). Possibly some of these lower-level ECSs are thereby driven out of their range of possible control and have also to reorganize. One can generate an avalanche
of reorganization this way, if the system as a whole had been operating near its limits; but in most cases, only one or a few ECSs are likely to be affected.

>If the change reduces intrinsic error, the rate of reorganization
>slows. When intrinsic error drops below some minimum amount, reorganization
>stops, and whatever organization of behavior exists at that moment has been
>"learned." It doesn't matter what that organization is, as long as it
>results in correction of intrinsic error.

The word "intrinsic" is what I object to in this. By making intrinsic error a stimulus and reorganization a response, you do two bad things--you bring into existence a separate reorganization mechanism, and you do not differentiate
among parts of the hierarchy that are working badly or well. I would agree that persistent intrinsic error should induce reorganization, but probably only affecting the ECSs that take the intrinsic reference signals. I would keep everything local, within the ECS: persistent error in an ECS induces reorganization, and greater or prolonged error induces restructuring. Both are blind.

If restructuring happens, then we must note that both the output and the input of an ECS typically connect to the same lower-level ECS (in an efficient hierarchy). Accordingly, not only do the actions affected by the higher-level ECS change, but so does the CEV that it perceives. This change can be radical, involving aspects of the world previous ignored completely.

Without restructuring or reorganization, the perceptual function of an ECS can change, either smoothly through a process such as Hebbian learning, or radically, as might happen if the perceptual function were programmatic (e.g. "if input A > input B, then report 1, else report 0 as the percept" might change to "if input A < input B, then report 1, else report 0"). Hebbian learning is guided learning. If an ECS has a percept akin to some useful CEV, then by Hebbian learning, it may come to perceive that CEV more precisely over time. In this case, Bill's comment:

>There is no advance specification of what is to be learned; anything will
>do, as long as it has the side-effect of reducing intrinsic error.

will not apply.

Another point on which I differ from Bill:
> I assume that the changes are all
>quite small, and that they occur at a low enough rate that the consequences
>of each change have time to be reflected in the state of the intrinsic
>error signal. I assume that a law of small effects is at work: small
>changes have small effects.

If the hierarchy incorporates a category level, the law of small effects will not hold. Small changes may have near zero effects most of the time, but very large effects some of the time. And when control at different levels occurs with different bandwidths, the rate at which the effect of any change can be assessed will vary at least as drastically, probably more so. The disturbances of a CEV by the world become very hard to distinguish from the effects of a change in the sign of one small component of a complex feedback system.

In sum: I buy the concept of random reorganization, but not that of a reorganizing system, or the idea that reorganization depends only on intrinsic error. And not all learning depends on reorganization (even if it incorporates what I call restructuring). Some learning is adaptive and directed toward regularities of the world.

Martin

Date: Tue Jun 30, 1992 11:04 pm PST
Subject: Reorganization

[From Bill Powers (920630.2000)]

Martin Taylor (920630.1500) --

Well, good. It was beginning to look as if we would soon run out of things to disagree about.

Let me acknowledge first several valid objections you have raised to my concept of reorganization. Yes, there is a problem with getting reorganization to occur in the right place, and not disturb systems that are working correctly (or at least usefully). Yes, there is a problem with the highest level of reference signals if they are not supplied by a reorganizing system. And (to a much lesser degree) reorganization of systems that handle discrete variables is a problem -- as everyone who has tried to build a self-writing program has discovered (no law of small effects there, as you point out, at least in a digital computer).

There is one way out of the first problem. That is to define the reorganizing system as a distributed system, a mode of operation of every ECS, but one that is NOT concerned with the normal business of control. This would solve the problem of specificity of the locus of reorganization, in that this distributed system would sense error and act to correct it at the place where it occurred. I have long held this concept in reserve -- I think I even mentioned it in BCP -- but as I don't have any idea what this special mode of operation might be (the Hebbian solution is not yet, to my mind, worked out well enough to model) I have elected to go with a lumped model that will work in essentially the same way. There are other possible solutions to this problem. And there are also reasons for NOT wanting too great a degree of specificity, and for NOT confining the motives for reorganization to purely local conditions, as will be seen as we go along.

The second problem, that of the highest level of reference signals, is not so difficult once the first problem is accepted as being solved (one way or another). The simplest solution is to say "I don't know" and wait for a better idea. Next would be to recognize that reference signals are, in themselves, meaningless; it's the perceptual function that gives a particular reference signal the significance of "so much of THIS perception." So a reference signal can simply be a bias in the perceptual function, or be set to zero (in which case the new perception will be avoided -- not such an unrealistic deduction!). Another possibility, under the hypothesis that reference signals are derived from memory, is that the most predominant experience of a particular perception at the highest level at any time during development becomes the default reference signal: this is simply the way the world works, and we defend it even not knowing why we do. That, too, does not seem so very unrealistic. And finally, we might suppose that the selection of the highest reference signals is always random, experimental. In that case we would have to allow the reorganizing system to select reference signals -- but of course not systematically, and as I will propose, not directly.

The third problem is the most difficult, and I think I'll pass on it for now because there are other aspects of reorganization that we need to discuss. I hope I'm not leaving out a problem that you consider more important than any mentioned so far.

On p. 188 of BCP there's a diagram of the relationship of the reorganizing system to the learned hierarchy, which I now wish I had drawn a little differently. The way it's drawn, one could easily see the reorganizing system as the highest level in the hierarchy, but for one rather subtle fact -- too subtle for many people to notice. If you look at the outputs of this highest-level system, you will see that they affect ALL levels in the hierarchy, not just the "next highest" level. That's not allowed in a hierarchical system, because intermediate levels of system will sense the result as a disturbance and alter their actions to counteract it. Levels can be skipped going upward, but not downward.

What I had in mind would have been better illustrated by drawing the sensors, comparator, and output function of the reorganizing system BESIDE the hierarchy rather than above it, with the output arrows travelling sideways to reach all parts of it. The main point of the diagram, however, is clear: the reorganizing system monitors variables that are neither sensed nor controlled by the neuromotor hierarchy. I'll try to justify that.

The need for some sort of learning system was always evident, even when the model was in its birth throes. Like everyone else, I assumed that a system responsible for growth and development would cause RELEVANT learning to take place. That is, if the system were hungry, it would learn how to get something to eat. If it were cold, it would find a way to get warm again. But the longer I tried to think of a way to make a system like this work, the more circular the problem seemed. You had to know that food would cure hunger before you could learn how to look for food. You'd have to know what food looks like and how it smells and tastes before you had ever seen, smelled, or tasted it.

Then I realized the cause of the difficulties: I was assuming that there was some natural obvious imperative relationship between feeling hungry and eating, between feeling cold and doing something like exercising to keep warm. I was assuming as a reason for learning the very thing that had to be learned.

I had read Ashby's _Design for a brain_ by this time, and (whether he thought of it this way or not), his concept of superstability, his "homeostat," showed the answer. The basis for reorganization can't be any PRODUCT of reorganization. It has to be something completely aside from WHAT is learned. Ashby called these somethings "critical variables." They are variables in the system that, if maintained within certain limits (Ashby's way of seeing it) or close to certain reference levels (my way), would assure or at least promote a viable organism.

It was then only a short step to realizing that if a hierarchy of control were ever to come into being, this process of reorganization had to be operational from birth, and most likely from early in gestation. This meant that it could not use any perceptions of intensities, sensations, configurations, transitions .... system concepts, as the nervous system would come to perceive such things, before the ability to perceive (and, I would now add, control) such things had been developed. This immediately ruled out any principle of reorganization that uses any familiar perceptions or means of control; particularly, any programs, principles, or system concepts. Those things would EMERGE from reorganization; they could not cause it. There could be no reorganizing algorithm.

What, then, could possibly be the basis for reorganization? What would guide it? Part of the answer lies in preorganization of the nervous system -- organization at least to the degree of making it possible to construct perceptual functions, comparators, and output functions, or their equivalents, with the necessary interconnections. But that only provides the possibility of an adult organization; the details must be left up to interactions with the environment, or no learning could happen, particularly not on the massive scale of learning that marks human development.

Well, I thought, why not pleasure and pain? When we feel pleasure, we feel no urge to change; when we feel pain, something is wrong and we must learn to behave differently, or at all. But pleasure and pain are abstractions, whereas a real organism must operate in terms of variables and processes relating them. Pleasure and pain are just labels for certain ill-understood experiences. What underlies them?

That led to the concept of intrinsic variables. By "intrinsic," I mean that these variables have nothing to do with anything going on outside the organism. They are concerned with the state of the organism itself. They might BE variables like blood pressure and CO_2 tension, or they might be signals, biochemical or even neural, standing for such variables. In general we can speak of them as signals, since the null case is that in which a variable IS the signal. The important feature of these signals is that they must be inheritable: they must exist in the organism prior to any reorganization. They are Ashby's critical variables.

For each intrinsic variable or signal, there is some state that is at the center of a range that is acceptable for life to continue. That is a _definition_ of an intrinsic variable and its reference level in this context, that separates such a variable from all the other variables in a living system. Ashby's term, critical variable, is probably better, and maybe after everyone understands what I'm talking about we might think about adopting it "officially." At any rate, for every intrinsic variable there is some reference level, so that when all intrinsic variables are at or near their respective reference levels, the organism is in a viable state.

Now the outlines of a reorganizing system begin to appear. When intrinsic variables depart from their normal reference levels, something is seriously wrong; survival is in question. This is "pain," generalized. If the organism is to survive, it must do something.

But in the beginning, it doesn't know how to do ANYTHING. It has no conception even of an external world. It has no idea of how that external world bears on its well-being. It has no knowledge of how to affect that external world even if it has the capacity to do so. Therefore, we have to conclude, whatever is done about the intrinsic error, it must be done at random -- without any systematic relationship to the outside world.

We can debate, of course, just how much of a head-start evolution actually provides for this process. In lower organisms, it's quite a large amount. But I wanted to solve the worst case because that would establish an important principle; any organization capable of working in the worst case would naturally work even better with a head start. So we can ignore that consideration.

In my model of the reorganizing system, I posit intrinsic reference signals and a comparator for each one. This is a metaphor; in fact, all we need assume is that there is some reference state established by inheritance, and that deviations of the intrinsic variables from their respective reference levels lead to reorganization. We don't need to guess at the mechanisms or even the locations of these processes -- that kind of question can be answered only by a kind of data that nobody knows how to obtain yet. We can only speak of functions, not about how they are carried out. A control-system diagram illustrates the functions and how they must be related.

The diagram of the reorganizing system on p. 188 in BCP shows just one intrinsic variable and perceptual signal, one reference level, one comparator, and one branching output that mediates the random changes in the target location, the present or future hierarchy of control systems. This is a schematic representation of a system that may involve hundreds of intrinsic variables with specific reference levels, and hundreds or thousands of pathways that connect error signals (if signals they be) to the target locations for reorganization. The signals may be purely biochemical, or some of them may be neural, although not part of the main hierarchy (the autonomic system and reticular formation may be involved). In all likelihood, this system that is shown as a single control system really consists of a multitude of control systems distributed throughout the body and nervous system, or throughout the biochemical systems which pervade everything. The geometry is immaterial, as is the nature of the signals and computers.

Now the crucial part: closing the loop for these intrinsic control systems that create the organization of behavior.

The outputs of the reorganizing system change neural connections, both as to existence and as to weights. They cause no neural signals in themselves; they merely change the properties of the neural nets. In doing so, they can connect sensory inputs to motor outputs, and thus, in the presence of stimulation, create a motor response to a sensory stimulus. They don't create any particular response; they only establish a functional connection so that motor output bears a relationship to sensory input according to the amount of input, should any such inputs occur.

The motor outputs affect the world, which affects the sensory inputs. The only stable configuration is that involving negative feedback. When there is negative feedback, some part of the world tends to be stabilized, or even to be brought into a specific state that is resistant to disturbance.

Of this negative-feedback (or other) relationship between action and perception, the reorganizing system knows nothing. The entire world of the reorganizing system consists of intrinsic variables, which relate to the state of the organism itself, not to the state of the outside world. But when actions occur, they affect the world, and the world affects the state of the organism in ways other than sensory. The state of the world affects intrinsic variables.

Therefore if reorganization results in stabilizing certain aspects of the external world against disturbance, and brings those aspects to specific states, the result may be -- MAY be -- to bring some intrinsic variables closer to their reference states. This is purely a side-effect of what the new control system is doing. What the control system is sensing and controlling may have nothing directly to do with the side-effect that is changing an intrinsic variable. But if the result of sensing and controlling in that way is to lessen intrinsic error, reorganization will slow or even cease. And that control system will continue in existence.

The question of specificity of reorganization arises. I hope you can see that the reorganizing system is loosely enough defined to allow for many possible solutions to this problem. I won't go into them just yet.

Now we can see the basic logic of reorganization. The reorganizing system is not concerned at all with what control systems exist or what variables in the outside world are brought under control. All it is concerned with is keeping intrinsic variables near their reference levels. If there is intrinsic error, reorganization commences, with the result that perceived variables are redefined, sensitivities and calibrations change, means of control change, and the external world is stabilized in a new state. The only significance of that new state to the reorganizing system is that intrinsic error may be corrected, putting an end, for a while, to reorganization.

Given a reorganizing system that works this way, an organism can learn to survive in environments having almost completely arbitrary properties. A pigeon with such a reorganizing system can come to maintain its internal nutritional state near an inherited reference level by pecking on keys instead of grain, or even by walking in figure-eight patterns -- it can learn to control variables that have absolutely no natural, logical, or rational connection to nutrition, and by doing do, can protect its own nutritional state. It does not need to reason out why performing such acts is so vital, or even what the connection is with getting food. It doesn't even have to know that ingesting little bits of brown stuff has the effect of keeping it from starving.

Reorganization is not an intelligent process; it produces intelligence as a byproduct of its real function, which is keeping the organism alive and functioning. It is the most powerful and general control system there can be, because it assumes nothing about the properties of the outside world. NOTHING. It does not even know there is an outside world.

If you recall my posts on the origins of life a year or so ago, you'll realize that this process of reorganization exemplifies exactly the same process that I proposed as the way the first living molecules came into being. The main difference is that the variability that creates new organizations to be retained or evolved away comes not from external forces but from an active process of random change driven by internal error signals. The reorganizing system is evolution internalized and made purposeful.

So it is important that reorganization not be TOO specific. In an arbitrary environment, there's no telling what control processes must be learned or modified in order to keep intrinsic error near zero. Reorganization depends on the ACTUAL effects of controlling certain objective variables, not on our symbolic understanding of experience, on our theories, or on our perceptual interpretations.

If there's one primary concept that must be understood to understand my theory of reorganization, it's that the variables controlled by this process are completely apart from the variables represented as perceptions in the learned hierarchy. The learned hierarchy is concerned primarily with sensory data about an outside world, and about those aspects of physiology that are represented in the sensory world. The reorganizing system is concerned about variables in the world beyond the senses -- with the actual state of the physical organism at levels inaccessible to the central nervous system.

Some of these variables might actually be in the brain. I have entertained the idea that because comparison is such a simple process, involving just a subtraction, comparators might be part of the kit with which we start building a functioning nervous system. In that case, error signals could also be intrinsic variables. It is not necessary to learn from experience that error signals represent some degree of failure to control; large error signals indicate serious problems, no matter what perception they relate to. The reorganizing system could monitor error signals in general, en masse, without any need to know what they mean, and their mere presence at large magnitudes could be sufficient to cause reorganization to start. This would satisfy the requirement that intrinsic variables be inheritable. And one result would be that loss of control could lead to highly localized reorganization, precisely in the system that has lost control.

But not all intrinsic variables are in the brain.

Another possibility is that certain kinds of intrinsic variables are associated with certain levels of control -- in other words, as Mark Olson suggested, that the levels of the hierarchy might be associated with different classes of intrinsic variables and reorganizing processes. This would at least localize reorganization at the right level, if not in the right system. But I have no idea what these classes of intrinsic variables might be -- what INTERNAL states would have special significance relative to the different levels of control and their effects in the OUTSIDE world. One might suppose that sexual intrinsic variables might require rather higher levels of control to be acquired and modified, as relationships with other control systems are involved. Achieving sexual satisfaction can certainly require walking in figure eights and worse. But examples are hard to think of.

I have also proposed that reorganization might be directed by awareness, and entail what we feel as volition. That's pie in the sky right now.

As to "restructuring," I don't feel that this necessarily relates to what I think of as reorganization. But I won't fight over this point, or over your proposals about the sequence in which various aspects of organization might come into being. The one point on which we apparently have a significant difference is on the relationship of intrinsic variables to what is learned. I hope I have explained more clearly what I mean by intrinsic variables and just why I think that there must be NO relationship to the learned control systems, save for the one imposed by the natural environment that relates the state of the world to the basic intrinsic state of the organism itself. It would seem that you have not understood my terms here; perhaps now you do.

Best, Bill P.

Date: Tue Jun 30, 1992 11:14 pm PST
Subject: DofF, alerting, etc (v. long)

[Martin Taylor 920701 0240] Today is Canada Day, the 125th and possibly last.

The following posting has been gestating for about 4 weeks. I hope it is intelligible.

Here is a very long statement (over 450 lines) of something that I think is an important consequence of Perceptual Control Theory, in a direction that has not been discussed, so far as I remember, in the last 18 months of CSG-L, or perhaps ever.

It deals with parallel and serial processing, similarity and difference detector systems, alerting systems and situation awareness, and finishes with a speculation about hyperlexia and autism. If you want to read it at all, you may prefer to print it out rather than to read it off the screen.

Warning: the latter part of this note is hand-waving speculation. But the earlier part represents what I think is necessarily true, given the basic premises of PCT.

Martin

Summary

The sensory systems provide many orders of magnitude more degrees of freedom for input than the skeletal system permits for output. Some reduction of the input degrees of freedom can be achieved by exploiting the natural redundancy of the world, but there remain many times more possible degrees of freedom for perception than there are for action. Hence, not all percepts that can be controlled are being controlled at any given moment. This discrepancy leads to the conclusion that there exist Elementary Control Systems (ECSs) whose perceptual input functions determine percepts that are not being controlled. These ECSs may be passive (only passing on the result of their perceptual input function) or monitoring (controlling, but through the imagination loop rather than through the real world).

Maintenance of percepts in desired stated is accomplished through the operation of active ECSs, but a monitoring ECS can determine when its percept is departing too far from its reference levels. If it is to achieve control, restoring its error to a tolerable level, it must wrest control from some active ECS, relegating that ECS to a passive or monitoring role. Alternatively, it can emit an alerting signal that causes a sibling ECS with less tolerance for error to take control[1]. A monitoring ECS contributes to "situation awareness," a previously elusive concept that makes sense in the context of PCT.

The shift of status of ECSs among monitoring, passive, and active states demands some kind of switching, either within an ECS (changing its gain function, for example) or in some separate system that can move control around within the hierarchy.

A monitoring ECS requires tolerance for amounts of error that would cause an active ECS to emit a substantial output signal. Likewise, a passive template-based alerting detector may exist, which brings some part of the external environment under control (shifting attention is one way of putting it). Such a detector would also be expected to have tolerance for error. Both seem to perform much as the human "similarity detection" system--parallel, fast, and tolerant of error, in contrast to the "difference detection" system which is slow, accurate, and seems to correspond with what one would expect of an active ECS.

The differences between monitoring and active ECSs seem to provide a natural reason for the two "tracks" of processing that I postulated in 1983 to account for the results of experiments in reading. In the process, they plausibly account for the relation between a pathology called "hyperlexia", and autism. Studies of overt control seem unlikely to be able to penetrate far into the functions of perception, since at any one time, most of perception is not under control.

Introduction

ln 1983 I published something I called the BLC (Bilateral Cooperative) theory of reading [2]. At the time, it seemed justified by a host of surrounding data, but lacked an evolutionary rationale. I now think that PCT gives it one, and in the process relates many concepts, including alerting functions, situation awareness, autism, and the role of attentional focus. The argument follows from the degrees-of-freedom (DOF) factors that I have mentioned on CSG-L from time to time over the last few months.

Background on perception in reading:

Logically, the concepts of similarity and of difference seem related, in that one is the polar opposite of the other. If two things are more similar, they are less different. Psychologically, that seems not to be so. In many respects, judgments of similarity act differently from judgments of difference, to the extent that they seem to produced by quite independent processes, and indeed to be preferentially processed by different brain hemispheres (I'm not up on the recent literature in this area, but this was the way it seemed in 1983, and I doubt that new data would much change this general statement).

Similarity processes seem to work fast, in parallel, and possibly unconsciously, whereas difference processes tend to be slower, serial, and attentive. Similarity processes tend to work on whole entities, whereas difference processes analyze the components of the entities. Where exactness of identification is required, the difference processes dominate ("Is this exactly an X?"), whereas if plausibility is more useful, the similarity processes work faster and may be used without the slower difference processes ("Might this be an X?"). Similarity processes can give multiple answers ("This is like an X, like a Y, and like a V"), whereas the difference processes can give only one answer in most situations ("This is not an X or a V, but is a Y").

The BLC model of reading postulates a layered set of levels of abstraction

for the perception of elements of text, ranging from the visual features of letters up to the concepts inherent in the material. At every level of abstraction, parallel similarity processes and serial difference processes are both available for the construction of the next level of abstraction. The degree to which each is used at any level depends on several factors, including the skill of the reader, the familiarity of the material, the importance of exactness, and so forth. A highly skilled reader skimming familiar material would tend to use largely the similarity processes at all levels, to form very quick impressions of the conceptual content of the material, whereas a poor reader might work through the same material syllable by syllable using the exacting difference processes to avoid error, but doing it very slowly. The skilled reader might also use the difference processes if reading for copy-editing, or as an editorial critic looking for logical flaws or rhetorical imprecision.
There is an interesting case of a French speed-reader who, after a stroke, lost the ability to read slowly, but could still speed-read [3]. At a very minimum, this case shows that the ability to analyze words slowly is not a prerequisite for reading them fast.

PCT background--degrees of freedom:

By my very rough count, the joints of the human body, together with such shape-changing functions such as facial expression and voice, admit around 125 degrees of freedom. There are probably far fewer, inasmuch as it is hard, for example, to flex the top finger joint while keeping the others straight, but I want to make this number as high as is reasonable. This number, 125, is the largest number of degrees of freedom that can be used in total by ANY set of ECSs within a layer of the hierarchy to control their actual perceptions from the environment. More ECSs in a layer may be involved in control, but then their perceptions will not be mutually independent[4]. Of course, the degrees of freedom controlled by the ECSs in one layer depend completely on those controlled by ECSs in the layer before, so the number of controlled degrees of freedom cannot increase (but may decrease) as we go up the layers, and the total number of independently controlled degrees of freedom overall cannot exceed about 125.

Again by a very rough count, the degrees of freedom for sensory input can be estimated as follows: about a million optic nerve fibres, about 60,000 auditory nerve fibres, many (I have no idea how many) touch and pain sensors on the skin, and quite a few taste and smell sensors. The actual numbers are not important; the point is that there are far more than there are degrees of freedom for affecting the external environment, by a factor of many thousand. Each incoming sensory nerve fibre IN PRINCIPLE (though not in practice) provides a degree of freedom for sensation. Let's take a low number for how many, and say 2.5 million, or about 20,000 times as many as there are controllable degrees of freedom. How can ECSs with 125 degrees of freedom for output control perceptions with 2.5 million degrees of freedom? The answer is that they can't, not all at the same time; but over time, the system could change which perceptual degrees of freedom are being controlled, so that any of them is potentially the subject of control.

To compound the problem, the intensity values of the sensory degrees of freedom can vary faster, usually much faster, than can the angles of the joints. It would probably be fair to put a limit of 1 to 2 Hz (cycles per second) for the average rate at which the joints can be oscillated, although some can go faster. The visual, touch, and auditory fibres can all vary their outputs much faster, though (to pre-empt myself) some of their speed comes from their working in a coordinated fashion. It would probably be conservative to allow the sensory fibres an average bandwidth of 10 Hz. Given these numbers, and assuming that the control of the joints can be no more precise than that of the sensory inputs, it is not unreasonable to think that the incoming information flow is in principle capable of reaching a rate about 5 orders of magnitude greater than can be controlled.

In practice, the world is not so unkind as to provide us with information that changes character in every part of the visual field 10 times per second. There is a great deal of redundancy. A chair remains a chair; neighbouring points tend to have much the same brighness and colour as each other and as they themselves did a second or two previously. Only at edges and when new objects come into view do the intensities change rapidly in space or time. But it is unreasonable to think that only one part in 100,000 of the incoming information is useful and new. Even if the redundancy is 99.9%, there is still 100 times as much information coming in as can be controlled for. And 99.9% seems very high. Regardless of the numbers, it is clear from immediate experience that there is a lot of sensory input for which we are not at any particular moment controlling. The point of the numbers is to suggest how very much uncontrolled perception there may be.

For two reasons, we can argue that there is no sensory input for which we inherently cannot control. The first is that it would violate the efficiency we expect of an evolved organism. Just as behaviour is the control of perception, so perception is the mechanism whereby we can act so as to survive long enough to propagate our genes. These actions control the relevant perception. Any perceptual capability that does not support this objective (cannot be controlled) will be unhelpful baggage, in an evolutionary sense, and will be selected against if it entails any cost.

The second argument comes from the theoretical arguments and experiments of J.G.Taylor, who showed at least how much more readily perceptual abilities are developed if they are the subject of control than if the corresponding sensory input is passively observed and identified by a teacher. His theory proposes that there is NO perceptual capability unless it can form part of a control loop, and this agrees with the evolutionary argument.

If there is no perceptual function that cannot be the subject of control, and if at any moment only a tiny proportion of the sensory degrees of freedom are under control, then (1) there must be some way of changing which perceptual degrees of freedom are under control from moment to moment, and (2) there must be some way that the living organism can detect which perceptual degrees of freedom should profitably be controlled at any moment. These two requirements are fundamental to the theoretical argument, and all of the foregoing has been devoted to showing that they almost necessarily are real requirements on a living system of the complexity of a human, not options.

ECSs that are not at a particular moment controlling their perceptions through the environment might nevertheless be controlling them through imagination, so they are not necessarily inactive as controllers. We can call them "impotent" or "monitoring" ECSs, since their operation does not affect any CEV in the real world at that moment. Other ECSs may not be controlling at all, having their output gain set to zero, while nevertheless continuing their perceptual function of abstracting from the incoming sensory data and passing the abstraction on to higher level ECSs. We can call these ECSs "passive." (Still others might be turned off completely, but we can ignore them in this discussion; all of the ECSs of interest have their perceptual input functions in operating normally).

Monitoring or passive ECSs can be supplied with reference signals, just as if they were actively controlling, and therefore can develop error signals that lead to outputs that provide references for lower ECSs. The only thing they cannot do is to have their outputs actively affect a CEV outside the person. If they control anything, it is through the imagination loops of ECSs in their part of the hierarchy. Through the imagination loops, they could control their prediction, or plan, for what they may perceive, but they cannot control what they actually perceive from sensory data. Most ECSs must be monitoring or passive at any given moment, but there is no reason to believe that their perceptual functions are switched off.

To recapitulate, there are, at any moment, three kinds of interesting condition in which an ECS might find itself: (1) "active," controlling its perception through a loop that uses muscular function; (2) "impotent" or "monitoring," controlling its prediction for incoming perception through loops which are completed only through imagination; (3) "passive," not controlling for anything, but nevertheless performing its perceptual function within the hierarchy.

Situation Awareness

"Situation awareness" is a nebulous concept that has often been associated with workload assessment. A naive reading of the words suggests that there is some kind of conscious awareness of the state of a complex environment. In PCT terms, consciousness has no explicit place, and, I believe, it should not have a place in assessing situation awareness. In everyday life, one is constantly acting within a complex and changing environment without being aware of acting appropriately in respect of disturbances in that environment. Operationally, one is "aware" of the disturbances without being conscious of them.

From the degrees of freedom argument, only a small proportion of the percepts based on the environment are being controlled, but it is possible that many more are being monitored through imagination-based control. Within HPCT, imagination-based control is like real control, except that at some level of the hierarchy the output is connected back to the perceptual input by an "imagination loop" rather than by the action of subordinate ECSs that eventually act through the real world. Hence, imagination works much like real control, with two important exceptions: (1) lower-level conflicts may not occur, permitting the simultaneous satisfaction of references that could not be simultaneously satisfied by control of the real world [5], and (2) imagination-based control can work very fast, not being constrained by the dynamics of objects in the real world. Imagination therefore can perform a predictive function, projecting what percepts would be obtained if the control were effective.

It is tautological that active ECSs (those really controlling percepts from the environment) are aware of the situation in respect of the percept that they control (if one removes the concept of consciousness from the connotations of "aware"). It also seems reasonable to suppose that monitoring ECSs are aware, in that they are shadowing the control of their percepts. The output of a monitoring ECS has no effect on the percept it receives, but the ECS can continually assess what effect it would have if it did take real control. A monitoring ECS can become passive simply by reducing its output gain to a negligible level, but it cannot become active without denying control to some other ECS. Passive ECSs (not providing any output to lower levels) do not need to make any use of their perceptual input such as to present it to the comparator, and cannot necessarily be considered to be situationally aware.

Of course, to identify situation awareness as being associated with monitoring rather than passive ECSs is speculation, but it is speculation with some justification. A monitoring ECS is controlling through imagination, and thus is prepared to take active control if necessary, without the introduction of an initialization transient. This ability seems to correspond with the notion of situation awareness, according to which the individual can react to the exigencies of a situation without the need for an initial period of updating the perception of that situation. According to this view, situation awareness is connected not so much with perception, but with the state of perceptual control implicit in monitoring as opposed to passive operation.

Alerting and parallel/sequential function:

ECSs can change state. A monitoring ECS can become an active controller, but if all the available output degrees of freedom had been used, some other ECS must change from the active state to being either a monitor or passive [but see again note 4]. Under what conditions should this happen?

One low-level example of an induced shift of control can be found in the motion-detecting system of the visual periphery. The central part of the visual field, imaged in the part of the retina called the fovea, is used for detailed, coloured vision. The periphery, on the other hand, has relatively poor resolution for detail, and the further out one goes, the less colour sensitivity exists. What does exist is a motion detection system that tends to lead the eyes to fixate on any location at which an unexpected movement has occurred. The movement is not an uncontrolled S-R event, since it can be suppressed, but in the absence of such suppression, it tends to be an automatic response as if it were a simple S-R event. The effect is to allow higher-level ECSs access to detailed information about the moving object that would allow them to determine whether there exists a conflict with any relevant reference value. The perpiheral motion detectors say "there might be something important here" and the detailed examination permits other ECSs to determine whether there is.

The peripheral motion detector system can be said to provide an alerting signal that (usually) leads to a change in the perceptual degrees of freedom that are under control, by changing both the portion of the environment from which detailed visual sensation is received and the variables in the environment that are the subject of attention. They induce a shift of attentional focus.

Two almost opposite conditions seem appropriate for the issuance of an alerting signal: (1) the perception coming into a monitoring ECS deviates sufficiently either from its imagination-driven prediction or from its reference, or (2) the perception in some passive ECS (or equivalently in a specialized pattern-matching system that is not an ECS) matches sufficiently closely some predetermined template that signals the need for control. A mother's awakening to the cry of her baby in a noisy environment might be an example of the latter.

Both proposed conditions for alerting signals depend on two things: rapid parallel operation of many possible alerting entities, and a loose tolerance for the identity of a perceived state with some reference. In the case of a monitoring ECS, the alerting state should occur only when the actual percept departs too far from the predicted percept (or perhaps when either departs too far from the reference signal in the ECS). In the case of the template-driven alerting signal, the fact that the relevant percept is not being controlled means that it will be highly variable and thus should be tested with a wide tolerance for error. Each of these conditions employ the concept of similarity rather than of identity.

I have been talking of "alerting signals" as if they were real signals that are sent from the ECS that detects the problem to some entity that has the responsibility of switching control among the myriad perceptual degrees of freedom. But this is not really necessary, especially in the case of monitoring ECSs. If an ECS has a gain function that stiffens with increasing error, then a monitoring ECS with a sufficiently high error could wrest control from some unknown other. If this generated high error in other ECSs, they also would wrest control, probably in some other direction, and eventually the system would stabilize in a condition in which all the ECSs were working in a relatively low-error regime. Such a system could not, however, maintain any of its percepts in a zero error condition, because it requires that there be a tolerance zone for error around which the dynamic gain is near zero. Without such a tolerance zone, all the monitoring and active ECSs would be trying to control, which they could not do because of the insufficiency of output degrees of freedom. There would necessarily be a great deal of conflict among them. Although there might be stability, that stability would not be achieved with zero error levels.

The conclusion seems inescapable (usually a warning sign!) that the system must incorporate some kind of switching mechanism. Either (1) an ECS can switch between high and low (zero?) gain mode, or (2) there are at least two types of ECS in sibling relationships (tolerant ones with zero gain near zero error and insistent ones with high gain near zero error) in which signals from one sibling can cause the other to assert control, or (3) an ECS can change its perceptual function from moment to moment, or (4) there are separate subsystems for alerting and for control, the controlling subsystem being able to reconfigure itself in response to signals from the alerting subsystem.

I am sure that there are other possibilities, but I see no escape from some kind of switching mechanism. A massively parallel set of passive or monitoring elements must in some way direct the operation of a set of active elements that among themselves are capable of controlling sequentially all the degrees of perceptual freedom monitored by the parallel elements.

Of the four listed possibilities, (2) and (4) are mutually compatible. There could well be two subsytems, one consisting of tolerant ECSs that normally act as a parallel monitoring system, the other consisting of insistent sibling ECSs most of which are inactive at any particular moment.

This is very close to the BLC hypothesis, except that the BLC model was posed in terms of a two-way flow of expectations and recognitions, in which the RIGHT track consisted of a multilevel set of parallel, tolerant, passive recognition units that provided goals for the LEFT track's sequential, exacting, slow units, and the LEFT track selected for correctness, pruning the RIGHT track's efflorescence. In both tracks, higher levels provided expectations to the lower, while lower levels provided results to the higher.

Hyperlexia and autism

The BLC model predicted various kinds of dyslexia and failure of understanding as depending on the failure of one or other of the tracks at some level of abstraction. One book reviewer pooh-poohed the whole idea, dismissing it with the comment that if the BLC model were true, one would expect there to exist people with right-hemisphere damage who would fail to see jokes, and would take the world literally. Such people do exist, with exactly the symptoms that the reviewer thought were so silly as to scuttle the model entirely.

Hyperlexia is one pathology that was not predicted by the BLC model, but that is consistent with it. A hyperlectic child is usually more able than a normal reader to extract words out of fragmentary or noisy representations, but is usually very poor at connecting the words identified into coherent structures, or to extract meaning from connected text. In [2], we speculated that the hyperlectic child had devoted too much of his or her resources to the RIGHT track, thus impoverishing the LEFT. We also noted that the literature suggested a high probability that a hyperlectic child would be autistic. It now seems that the PCT approach, using the arguments persented above, accounts at least plausibly for this connection.

In a pre-PCT evolutionary argument supporting the BLC model, I suggested that the need for a LEFT track derives from a requirement to choose between conflicting overt actions, whereas the RIGHT track permitted a wide-ranging appreciation of the state of the world, having only a slender connection with current overt action. One of the jobs of the RIGHT track was to alert the LEFT track to possibilities for action, among which the LEFT track could make precise judgments and initiate the overt actions. These characteristics now seem to be those of active ECSs in the LEFT track, and monitoring (or passive) ECSs in the RIGHT.

If we were correct that the hyperlexic child devoted too many resources to the RIGHT track at the expense of the LEFT, it would now seem to follow that the child would have few active ECSs, and would not control many of its percepts. It would passively observe the world much of the time, until jolted by necessity (or perhaps by chance) into producing some overt action. It would be situationally aware, but not visibly active.

Of course, a child that interacted very little with the world would have little chance to develop effective ECSs, so the situation would feed upon itself. The imbalance between active and monitoring ECSs would tend to grow. The child has "learned" that active control does not work very well, and devotes resources instead to monitoring, so that control can be exerted in those finely determined conditions where it is really necessary. But there would be very little development of subtle discriminations, the learning of which would normally be based on continued interaction with the world, interaction not experienced by the autistic child.

Comment

Once again, PCT ties together many threads that were disparate, though possibly seeming to be connected for non-obvious reasons. I find that this argument makes the BLC model natural and plausible, whereas before it was just a description that seemed to fit the data of experiments on reading and related studies. The concept of "situation awareness" becomes less vague, though remaining hard to experiment with. Studies of overt control seem unlikely to be able to penetrate far into the functions of perception, since at any one time, most of perception is not under control.

----

[1] A "sibling" ECS means any ECS or set of ECSs that can control much the same percept as the one in question. The notion is meant to take into account the idea that there may be separate but related ECSs for monitoring and for controlling actively a particular percept.

[2] Taylor, I. and Taylor M. M. The Psychology of Reading. Academic Press, 1983. (The BLC material is mainly concentrated in Chapter 11, with related experimental data in Chapter 10)

[3] Andreewsky, E. DeLoche, G. and Kossanyi, P. Analogies between speed-reading and deep dyslexia: Toward a procedural understanding of reading. In M. Coltheart, K. E. Patterson and J. C. Marshall, (Eds.) Deep Dyslexia. London: Routledge and Kegan Paul, 1980.

[4] It is simplistic to say that one degree of freedom is controlled by one ECS, for various reasons. Nevertheless it is convenient in the following discussion to maintain this fiction, which does not affect the argument.

[5] Which may explain why people sometimes prefer to "live in a fantasy."

Martin Taylor
DCIEM, Box 2000, North York, Ontario, Canada M3M 3B9
(416) 635 2048